Gasoline Price Dispersion and Consumer Search: Evidence from a Natural Experiment

Michael D. Noel*

April 28, 2018

Abstract

The vast majority of empirical studies examining the link between consumer search and price dispersion focus on how changes in consumer search impact price dispersion. This article does the reverse – it examines how a shock to price dispersion impacts consumer search. A direct measure of search is used and an exogenous shock to price dispersion is found in a refinery fire that caused decades-old retail gasoline price cycles, and the non-linear high-frequency price dispersion pattern generated by them, to stop. Identifying effects from this shock, the results show a substantial response of consumer search to changes in price dispersion.

JEL Classification Codes: L11, L15, L81, L91, Q31

Keywords: Price Dispersion, Search, Search Costs, Retail, Gasoline, Edgeworth Price Cycles.

I Introduction

Economists have long studied the causes and consequences of consumer search and for good reason – understanding how and why consumers search has important implications for competition, prices, profits and consumer welfare. Stigler [1961]'s seminal result is essentially that consumers equate the marginal benefit of search to the marginal cost of search and, since that time, a large theoretical literature has examined search in an equilibrium setting. One finding that repeatedly surfaces is that search activity and price dispersion are jointly determined. In competitive environments, greater consumer search tends to lead to lower levels of price dispersion (since high priced firms would be more severely punished) while lower price dispersion tends to lead to lower levels of consumer search (by reducing the benefits to searching).

*Author's Affiliation: Department of Economics, Texas Tech University, Lubbock, Texas, U.S.A. *email: michael.noel@ttu.edu.*

A large empirical literature has followed which tends to examine the price-dispersion-consumer-search relationship only in one direction – how exogenous shocks to consumer search affect equilibrium price dispersion. The literature is largely silent on the interesting but reverse relationship – how shocks to equilibrium price dispersion affect consumer search. The omission is surprising because the two are jointly determined and price dispersion is a key determinant of the benefits of search, and therefore search. In this article, I contribute to the literature by examining this largely unstudied reverse relationship.

The limited attention stems from two practical difficulties. First, a plausibly exogenous shock to equilibrium price dispersion is hard to find and, second, a direct measure of consumer search is hard to find. In this article, I overcome both obstacles in an examination of the retail gasoline industry in the province of Ontario, Canada.

The exogenous shock to price dispersion comes in the form of a rare natural experiment – a refinery fire that caused decades-old retail gasoline price cycles (known as Edgeworth price cycles) to suddenly stop in several nearby cities. As is well-known from the literature, high-frequency Edgeworth price cycles create a distinct, non-linear, and high-frequency pattern in price dispersion over the course of a cycle, oscillating between extremely high and very low levels, in synchronicity with the phases of the cycle. When the cycles in these cities stopped, the cycle-induced pattern of price dispersion also stopped. This is not to say that all price dispersion stopped, of course – there will be still price dispersion even in the absence of the cycles – only that the non-linear high-frequency pattern generated by and synchronized to the cycles would be no more.

As background, Edgeworth price cycles necessarily generate this pattern by their construction. Price dispersion levels spike to exceptionally high levels during the brief period when prices are quickly rising, as firms stagger large price increases in what is known as the 'relenting phase', and then suddenly collapse to very low levels when prices start to fall, as firms undercut one another by small amounts in what is known as the 'undercutting phase'. While price dispersion levels tend to rise some as the undercutting phase progresses and as firms undercut each other at different rates in different neighborhoods, they remain well below relenting phase levels. The result is a non-linear, high-frequency price dispersion pattern generated by, and synchronized to, the phases of the cycle – from extremely high price dispersion in the relenting phases to very low or low price dispersion in the undercutting phases. The universal relationship between retail price cycles and price dispersion is a theoretical necessity and has been confirmed empirically where it has been tested.

In the markets studied here, the retail price cycles were daily price cycles prior to the fire (i.e. they had a period of one day), meaning that the high-frequency pattern of price dispersion would repeat itself every day. The large price dispersion spikes, and very high levels of price dispersion, would occur over a few hours in the mornings when firms were increasing prices in a staggered fashion, by upwards of 20-25 U.S. cents per gallon. Substantially lower levels of price dispersion would dominate the rest of the day when firms were tit-for-tat undercutting by only a few pennies at a time.

In contrast, when the retail price cycles stopped after the refinery fire and the large spikes in price dispersion were no more, the degree of intraday price dispersion would almost certainly become more uniform over time, i.e. relatively more consistent, from hour to hour over the course of a day. There would still be price dispersion and potentially some intraday variation due to other factors, but without the strong and overriding pattern imposed by the cycles, there is no longer any clear reason for price dispersion to spike to extremely high levels for a few hours each morning and then suddenly collapse to very low levels the rest of the day, day after day. Consistent with the experience of other gasoline markets without daily retail price cycles, the result should be a more uniform, i.e. more consistent, intraday pattern of price dispersion. In fact, in markets without cycles, it is common for prices to remain unchanged throughout a day, or for several days at a time. Regular and predictable high-frequency intraday swings in price dispersion of the magnitude and kind experienced with retail price cycles has not been documented in markets without retail price cycles.

This sets up a clean natural experiment to study the effects of price dispersion on consumer search. Before the fire and in the presence of cycles, there is an exceptionally large difference in price dispersion between the early morning periods (relenting phases) and the midday periods (the start of the undercutting phases). After the fire, when price cycles ceased, I expect from the previous literature that price dispersion will become more uniform over the day. The fire thus acts as an exogenous shock to equilibrium price dispersion that can be used to examine search effects. By comparing consumer search in the early mornings to consumer search in the midday periods in particular, first before the fire and then after the fire, and then taking the difference in the differences, I can identify the effect of the differential intraday shock to price dispersion on the amount of consumer search, all within a single-city-contained difference-in-differences framework.

To my knowledge, this is the first study to investigate the effects of price dispersion on consumer search in a natural experiment setting. Exogenous shocks to price dispersion are difficult to find, and as equilibrium shifts between cycle and non-cycle regimes are exceedingly rare, the opportunity to use such a shock to equilibrium price dispersion to examine its effects on consumer search is unique. Chandra and Tappata [2011] comment on the difficulty of studying the price-dispersion-consumer-search relationship, stating that 'it is important to find a control group or benchmark' for identifying the relationship. The post-fire experience in the affected cities is a unique example of an almost ideal control group. I compare consumer search by the same consumers buying from the same firms in the same cities, all within days and weeks of each other, first experiencing regular and large temporary spikes in price dispersion at specific times of the day and then, just days later, not.

The second practical difficulty in investigating the reverse relationship between price dispersion and consumer search is finding a direct measure of consumer search. Fortunately, unlike many other brick and mortar industries, a direct measure of search is available for the retail gasoline industry. I use the number of price reports submitted to GasBuddy.com, a popular crowd-sourced gasoline price information website, by consumers who search for gasoline prices and then report the results of their searches to the website to help others in their own price searches. The measure is relatively new to the literature and, as discussed below, has some important advantages over other measures in this context.

To preview results, I find evidence of a significant and positive reverse relationship – higher equilibrium price dispersion leads to higher consumer search. At exactly those times of the day when price dispersion fell disproportionately after the fire in the affected cities – the early morning periods – I find that consumer search disproportionately fell as well.

I offer several contributions. The first is to empirically establish that an exogenous shock to price dispersion impacts the amount of consumer search, and to do so in a natural experiment setting. The second is to flag the potential for bias in other studies that examine how consumer search affects price dispersion (i.e. the usual direction studied) but do not account for the reverse endogeneity. Price dispersion and consumer search are jointly determined, and by understanding and accounting for the reverse relationship, it is possible to help quantify and correct the bias in those studies.

II Literature and Background

I draw upon two branches of the literature – the consumer search literature and the Edgeworth price cycles literature.

II.1 The Search Literature

There is a large theoretical literature on equilibrium consumer search in which price dispersion plays an important role. In a seminal article, Diamond [1971] finds that if all consumers have positive search costs, price dispersion falls to zero, prices rise to monopoly levels and, as the benefits to search are erased, there is no search in equilibrium. Varian [1980] and Stahl [1989] develop models in which some consumers have zero or negative search costs instead, and in these models price dispersion and search do occur in equilibrium. The theoretical literature shows that search and price dispersion are jointly determined and that, in competitive markets, greater consumer search leads to lower levels of price dispersion and lower price dispersion leads to lower levels of consumer search, all else equal.

Numerous empirical studies have followed and tested how shocks to consumer search affect price dispersion. Direct measures of search are rarely available, so most studies use a proxy for consumer search on the right hand side and a direct measure of price dispersion on the left hand side, with search proxies most commonly based on variables correlated with search costs or search benefits. The identification strategy in these studies is based on the idea that a change in the observable proxy will change either search costs or search benefits, and therefore unobserved consumer search in the background, in the expected ways.

Sorensen [2000] compares price dispersion in pharmaceuticals for acute illnesses versus those for chronic illnesses and finds that price dispersion is lower for the latter (where repeated purchases make the benefits of search, and presumably search, greater). Brynjofsson and Smith [2000] find that price dispersion is lower for goods sold on-line where search costs are likely to be lower than in brick and mortar stores. Other studies examining the effect of consumer search on price dispersion include Dahlby and West [1986] on auto insurance, Walsh and Whelan [1999], Zhao [2006], Dubois and Perrone [2015], and Sherman and Weiss [2017] on groceries, Baye et al. [2004] and Tang et al. [2010] for on-line purchases via shopbots, Milyo and Waldfogel [1999] on liquor, Orlov [2011] on airlines, and many others. For the retail gasoline industry, Barron et al. [2004] find that price dispersion is lower when stations are more densely located (which should lower consumer search costs), while Lewis [2008] finds that price dispersion is lower when there are more competing stations of a similar brand type. Tappata and Chandra [2011] find that price dispersion is lower across stations at the same intersection, while Pennerstorfer et al. [2015] find that price dispersion depends on the percentage of (generally better informed) commuter traffic present in an area.[1]

Notably all of the above mentioned papers estimate how consumer search (the right hand side variable) affects equilibrium price dispersion (the left hand side variable). The reverse relationship – how a shock to equilibrium price dispersion affects consumer search – remains largely unexplored.

Among the first studies of retail gasoline markets that use a direct measure of consumer search as a left hand side variable is Lewis and Marvel [2011]. While the study does not consider price dispersion, it measures consumer search using pageviews (or 'hits') to the GasBuddy.com website, and finds that search activity responds more to generally increasing prices than to generally decreasing prices over time. The result is consistent with the reference price model of Lewis [2011] and can potentially explain the well-known 'rockets and feathers' effect common in retail gasoline markets (Borenstein et al. [1997] and many others).[2]

Among the first studies to examine how price dispersion affects consumer search in retail gasoline markets, and the two most closely related to this one, Byrne et al. [2015] and Byrne and de Roos

[1] Hosken et al. [2008], in a study of the Washington D.C. area, opine that the degree of price dispersion across stations is relatively high considering the homogeneity of the product. In a related branch of literature, Nishida and Remer [forthcoming] take a more structural approach and seek to back out estimated search cost distributions from prices distributions directly, based on the search model of Burdett and Judd [1983] and the method of Hong and Shum [2006].

[2] Noel [2009] finds that asymmetric price cycles alone can result in a finding of rockets and feathers, even when the cycle taken as a whole moves symmetrically in response to price increases and decreases. Comparing cities with and without cycles, Lewis and Noel [2011] find that markets with asymmetric price cycles in the U.S. are also more price-fluid markets and lead to less asymmetry in passthrough overall.

[2015] use a direct measure of consumer search on the left hand side and find a positive relationship in OLS regressions of consumer search on price dispersion. The identification in these studies partly comes from the fact that the study markets in question – cities in Ontario, Canada in the former and the city of Perth, Australia in the latter – are characterized by high-frequency retail gasoline price cycles known as Edgeworth price cycles. As noted, Edgeworth price cycles generate a distinct, non-linear and high-frequency pattern of price dispersion, that to a large extent can be considered pre-determined.

A limitation of these studies, however, is that the price cycles they examine tend to be synchronized to a calendar week, each about a week long, meaning that the cycle phases tend to be collinear with daily changes in demand. Absent a shock to the presence of the price cycles themselves, it can be difficult to distinguish between two effects: a) changes in consumer search that are caused by price dispersion changes along the weekly price cycle, and b) changes in consumer search that naturally arise due to normal weekly traffic patterns, which follows a cycle of its own.[3] In this article, I build on these studies by exploiting a natural experiment that 'shuts off' the cycles and separates these two effects, allowing me to isolate just the price-dispersion-consumer-search relationship of interest.

II.2 The Edgeworth Price Cycle Literature

At the heart of the natural experiment is the sudden disappearance of decades-old retail gasoline price cycles in several cities surrounding the city of Nanticoke, Ontario, after a refinery fire there temporarily slowed production. In the following two sections, I review the literature on retail gasoline price cycles, what they are, and how they generate their unique price dispersion pattern.

High-frequency retail gasoline price cycles have been documented in numerous countries over the past few decades – e.g. Canada (Eckert [2003], Noel [2007a]), the United States (Castanias and Johnson [1993], Lewis [2009], Lewis and Noel [2011]), Australia (Wang [2009b], de Roos and Katayama [2013], Noel and Chu [2015]), Norway (Foros and Steen [2013]), and Germany (Allen

[3] For example, in Perth at the time of the study, price dispersion was highest on Thursdays as a direct result of the cycle (relenting phases occurred on Thursdays), but midweek is also the time when commuter traffic and consumer search would be high. Price dispersion was lowest on weekends with the cycle (price dispersion is lowest during the early part of undercutting phases), but this is also a time when commuter traffic and consumer search would be low.

et al. [2014], Siekmann [2017]). In many cases, the cycles have a period of almost exactly a week or almost exactly a day.[4] Figure 1 shows a picture of a retail gasoline price cycle in the city of Toronto, one of the cities examined in this study, showing a period of almost exactly a day.

The leading theory behind the price cycles is the Edgeworth price cycle model of Maskin and Tirole [1988], as extended by Eckert [2003] and Noel [2008]. In the original model, two identical infinitely-lived firms produce homogeneous goods and set prices in an alternating fashion using Markov strategies. Maskin and Tirole show that one of two possible types of equilibria can result – a focal price equilibrium (i.e. a constant price or constant markup) or an Edgeworth price cycle. Either type of equilibrium can occur under most supply and demand conditions and, once reached, tends to be persistent. Figure 2 shows an example of price paths in a theoretical price cycle, with marginal costs equal to zero.

In an Edgeworth price cycle, firms repeatedly undercut each other's prices by small amounts to steal market share back and forth. When prices reach marginal cost, a war of attrition ensues until one firm finally 'relents' and raises its price, all at once, to a much higher level. Others follow relatively quickly, and from the top of the cycle, undercutting begins again.

A large empirical literature has followed and produced some general results. Cycles are more likely to occur in larger markets (Noel [2007a], Valadkini [2013]), when there are more price aggressive firms (Eckert [2003], Noel [2007a]) and where station-level price elasticities are very high (Wang [2009a]).[5] Large firms tend to increase prices first and small firms tend to decrease prices first (Noel [2007b], Lewis [2012]), while mom-and-pop type stations are least likely to follow the price cycle (Doyle et al. [2010]). Numerous studies find that retail price cycles lead to lower prices (Noel [2002], Doyle et al. [2010], Zimmerman et al. [2013], Noel [2015], and Siekmann [2017]), while de Roos and Katayama [2013] and Foros and Steen [2013] express collusive concerns relating to the relenting phases. Lewis [2009] and Lewis and Noel [2011] show that markets with Edgeworth price cycles tend to recover from temporary cost shocks more quickly, while Noel [2012] and Noel and Chu [2015] show that price sensitive consumers can systematically purchase gasoline below the

[4] Cycle peaks synchronized to the day of the week or time of day are common (Noel [2015], Noel and Chu [2015], Foros and Steen [2013], Atkinson et al. [2014], Valadkini [2015], Siekmann [2017]).

[5] Using daily data for a set of cities in Ontario, Canada, Byrne [2015] finds that price cycles were most common in medium sized cities and less common in large cities. However, at that time, most large cities that had cycles had higher-frequency daily cycles that would not be visible in his daily price data.

average overall price by using simple rules of thumb that time the troughs of the cycle.

II.3 The Universal Relationship Between Retail Price Cycles and Price Dispersion

Central to my identification strategy is the relationship between retail price cycles and price dispersion. Station-specific data is not available for this study and I cannot measure price dispersion directly, so I must rely on this relationship to infer the pattern of price dispersion that existed over the path of the cycle prior to the fire. Fortunately, the relationship between retail price cycles and price dispersion is a universal one that is a necessary implication of the theory and has been confirmed empirically wherever it has been tested, including in the same study markets I examine here at slightly different times.

In short, in a retail price cycle, price dispersion suddenly spikes to very high levels during the brief time when prices are rising (during the relenting phases of the cycle) only to collapse to very low levels when prices are falling (during the undercutting phase), even lower in the earlier part of the undercutting phase than in the later part. The pattern repeats over and over in synchronicity with the phases of the cycle, according to the following mechanism. In a relenting phase, which may be as short as a few hours, firms increase prices by very large amounts and all at once at a given station, but at different times for different firms and different stations, i.e. price increases are staggered. This necessarily results in prices that are vastly different from one another for a short time – as much as 30% higher at some stations than at others (Eckert and West [2004], Noel [2007b], Doyle et al. [2010], Noel [2012]). (Noel [2015], Noel and Chu [2015]) – and all in an environment when price differences are otherwise generally very small. The gap between high and low priced stations necessarily leads to very high levels of price dispersion. Once all firms have relented to a similar price level, the undercutting phase begins in which firms repeatedly undercut one another by generally only a few pennies at a time. The price distribution collapses and price dispersion quickly falls to relatively very low levels. Price dispersion can then trend up over time as firms undercut each other at slightly different rates in different neighborhoods depending on local conditions (Eckert and West [2004], Bloch and Wills-Johnson [2010], Lewis [2012]), before it spikes up again during the next relenting phase.

The resulting within-cycle variation in price dispersion, synchronized to the phases of the cycle, is high. It is a necessary implication of the Edgeworth price cycle theory and the cycles would not look the way they do without it. The general pattern is easily seen in Figure 2, above, showing an example of the theoretical cycles. Price dispersion is very low in an undercutting phase – with a minimax range of just one unit on the price grid until, right at the bottom, it increases to two – yet exceptionally high in the middle of the relenting phase – with a minmax range of ten, or approximately ten times its average value in an undercutting phase.

The relationship has also been confirmed empirically. In real world relenting phases, the size of the price increase, and the consequent degree of price dispersion, is exceptionally high. The average relenting phase price increases in 2017 were approximately 15 to 20 U.S. cents per gallon in the Midwest U.S., 24 U.S. cents per gallon in Canadian cities, and 74 U.S. cents per gallon in major Australian cities.[6] That 74 cent per gallon difference is 30% of the average US$2.50 trough price in Australia, an extraordinary difference between neighboring stations. In contrast, during undercutting phases, stations are most commonly within a few pennies of one another and price dispersion is very low in comparison.

The relationship is so integral to the existence of the cycles that few studies explicitly mention it but those that do confirm the relationship. Byrne et al. [2015] show that price dispersion in a set of Ontario cities (as measured by the interquartile range) is approximately *three times* greater in the relenting phase than in the middle of the undercutting phase, a statistically significant difference.[7] Lewis [2012] shows that the range between bottom-of-the-cycle and top-of-the-cycle prices for a particular U.S. retailer is approximately *four times* greater the typical range for that same retailer in the middle of the undercutting phase.[8] Noel [2007b] finds that the median pairwise difference in the prices of any two stations in the city of Toronto in 2001 was just 0.4 cents per liter during the undercutting phase (1.0 U.S. cents per gallon), but the average price difference between stations at the 'top' versus the 'bottom' of the cycle in the relenting phase was a substantially higher 5.7 cents per liter (14.0 U.S. cents per gallon). The relationship is also clearly evident in the few studies that explicitly show station-specific graphs of the retail price cycles (e.g. Noel [2012]).

[6] GasBuddy.com (for the U.S. and Canada) and accc.gov.au (for Australia). Last accessed November 1, 2017.
[7] See Byrne et al. [2015], Figure 3.
[8] See Lewis [2012], Figure 3.

While the large difference between relenting phase and undercutting phase levels of price dispersion are most notable, the literature also shows that price dispersion can trend upward as the undercutting phase progresses, remaining below relenting phase levels. Byrne et al. [2015] shows that price dispersion at the end of the undercutting phase (on the last day before the relenting phase begins) is significantly higher than price dispersion at the beginning of the undercutting phase (the day after the relenting phase is complete), while Lewis [2012] shows a similar result for the U.S.[9] In each case, undercutting phase price dispersion levels continue to fall well short of relenting phase levels.

Because station-specific price data is not available around the time of the 2007 fire, I must assume that this relationship continues to apply to the markets I study here as it does elsewhere. Adding to my confidence is the fact that the relationship has also been explicitly confirmed for the same areas that I study here but at slightly different times, once shortly before my sample period and once shortly after. Byrne et al. [2015] examine numerous cities with cycles in the same province of Ontario, including cities that surround the study markets on all sides, but for a sample period a year later, and explicitly show that the universal relationship holds in these cities at that time. On the other end, Noel [2007b] examines the city of Toronto, one of the same cities examined here, except for a sample period a few years earlier, and explicitly confirms that the universal relationship holds in Toronto at that time as well. As the two studies confirm the universal price-cycle-price-dispersion relationship in these areas both before and after my sample period, essentially bookending the current study in time, it is highly unlikely that Toronto or London would have been anomalous in the interim or the first known and temporary exceptions to the universal rule. Nonetheless, absent station-specific data for the precise time period in question, this is an assumption that I make.

[9] See Byrne et al. [2015], Figure 3 and Lewis [2012], Figure 3.

III Data and Methodology

III.1 A Plausibly Exogenous Shock to Equilibrium Price Dispersion

The plausibly exogenous shock to price dispersion comes in the form of a refinery fire in Nanticoke, Ontario, Canada, on the shores of Lake Erie, on February 15, 2007. The fire caused a temporary supply shortage in several southern Ontario markets which in turn caused decades-old retail gasoline price cycles in several nearby cities, including the cities of Toronto and London (Ontario), to suddenly stop (Atkinson et al. [2014], Noel [2015]). When the cycles stopped, the non-linear, high-frequency pattern of price dispersion universally generated by the cycles – with price dispersion fluctuating between extremely high levels for a few hours during the morning relenting phases and relatively very low levels at other times, in synchronicity with the phases of the cycle – also came to a stop. There would still be some price dispersion after the fire of course, but the distinct pattern imposed by the cycles would be no more.

The Nanticoke refinery was one of several in the area supplying the populous Southern Ontario region. The fire was quickly extinguished after it broke out, and the refinery returned to full operation shortly thereafter. There was little press coverage of the event beyond a press release by the owner of the refinery to announce that the fire occurred and a second to announce its return to full capacity. Few stations ran short on gasoline, and beyond a temporary (and not historically unusual) price increase after the fire, there was little indication of a problem on the street. Continued production at the Nanticoke refinery and other refineries, imports from the U.S., and existing inventories were sufficient to meet demand. The most notable impact of the fire was that it caused the retail gasoline price cycles that had been present in several nearby cities for decades to suddenly stop.

That the refinery fire caused the stoppage of the cycles is consistent with the predictions of the Edgeworth price cycle model. Noel [2008] shows that when firms are sufficiently capacity constrained and cannot handle a sudden increase in market share, the incentive to undercut is reduced and the cycles necessarily stop. The non-reappearance of the cycles after the refinery returned to normal operation is also consistent with the model since equilibria tend to be persistent. Recall that the model is a multiple equilibria model – Edgeworth price cycle equilibria and focal price equilibria are both possible and persistent equilibria in the absence of capacity constraints (Maskin and Tirole [1988], Noel [2008]), and with no theoretical requirement to revert back to a cyclical equilibrium, stable prices remained in place in these cities for some time.

The primary study markets in this article are Toronto and London, both in the province of Ontario and both served in part by the Nanticoke refinery. Both experienced daily retail price cycles for decades before the fire (Noel [2007a]) and, in both cases, the cycles suddenly stopped after the fire.[10] I collect data on average regular grade gasoline prices reported to the gasoline price information website, GasBuddy.com, and the number of such reports, for each city within each of four intraday periods: 6 a.m.-10 a.m., 10 a.m.-2 p.m., 2 p.m.-6 p.m., and 6 p.m. to midnight, spanning from the morning of January 15, 2007 to the evening of March 14, 2007. I discuss the GasBuddy.com data source in more detail below. The collection and use of high-frequency intraday price data is novel – few studies use it (Noel [2007b], Atkinson [2008], Atkinson et al. [2014], Noel [2015], Siekmann [2017]) – but it is necessary to examine price cycles at this high a frequency. I supplement the data with a longer series of price and rack price data collected at a daily level, from February 15, 2006 to March 14, 2007, one year on either side of the fire. Summary statistics for the daily dataset and the four-times-a-day dataset are reported in Table I.

Figure 1, above, shows a close-up of the daily price cycles in Toronto in the weeks leading up to the fire. Toronto experienced daily cycles prior to the fire with relenting phases that regularly occurred in the early morning periods, before 10 a.m. It was similar in London. Figure 3 plots retail prices in each city in the weeks surrounding the February 15th fire and clearly shows how the retail gasoline price cycles disappeared within days of the fire. Figure 4 plots price-cost margins (using the rack price as a proxy for cost) over a longer period and shows that margins in each city became more stable after the fire as retail prices switched to a more cost-based pricing pattern. In the city of Toronto in particular, prices became almost perfectly correlated with rack prices over time, yielding an almost constant average five cent per liter margin for stations.

This sets up a clean natural experiment for examining price dispersion and consumer search. Prior to the fire, the universal relationship between retail price cycles and price dispersion means

[10] It was also apparent that the cycles stopped in the city of Ottawa, but that city was excluded due to insufficient numbers of price reports, 21 in total per day (less than four per station per month). In contrast, the total number of price reports was 187 per day in London and 1404 per day in Toronto.

that there would have been a strong and synchronized pattern of price dispersion over the course of a single day – extremely high price dispersion in the mornings, then very low, very low, and low, across the remaining three intraday periods. After the fire and in the absence of cycles, the universal relationship no longer applies and I appeal to the experience of gasoline markets that do not have cycles to understand intraday price dispersion patterns. In these markets, gasoline prices still exhibit price dispersion, of course, and the overall average level of price dispersion could be either higher or lower than in markets that have cycles (averaged across the cycle phases), but the daily pattern of price dispersion in these markets tends to be very different. Absent any clear reason for price dispersion to suddenly spike to very high levels for a few hours each day and then collapse to very low levels the rest of the day, day after day, price dispersion in markets without daily retail price cycles tends to be relatively more uniform over the course of the day, i.e. more consistent, from morning to midday to evening. In fact, it is common in these markets for prices at a given station to remain unchanged throughout the day or even for several days at a time, resulting in more consistent levels on an intraday basis. So for this study, I simply assume that price dispersion in the study markets after the fire would also become relatively more uniform over the course of a day, as expected. In other words, I assume that the intraday pattern of price dispersion in the study markets after the fire should qualitatively look like any other gasoline market in Canada or the United States that does not experience retail price cycles.

This is hinted at fairly strongly in Figures 3 and 4, which shows that prices and margins in the study markets became substantially less volatile after the fire. While the figures are based on city averages, it would be difficult to imagine how, on one hand, intraday price dispersion at the station-level would become even more volatile after the fire – with even larger spikes in price dispersion (only for a few hours each day and always in the mornings) – while on the other hand, Figures 3 and 4 clearly show that city-average prices and margins suddenly became substantially more mean-stable.

The differential intraday shock to price dispersion lends itself well to a difference-in-differences framework for studying consumer search, with a self-contained difference-in-differences estimation for each study market. The first difference is before versus after the fire, the second difference is the intraday period. The early morning intraday period, 6 a.m. to 10 a.m., serves as the primary treatment group, when price dispersion would have disproportionately fallen relative to the midday periods and when consumer search is also expected to disproportionately fall. The minor treatment group is the evening period, from 6 p.m. to 12 a.m., which received a more muted treatment shock that may result in lower search, relative to the change in search in the midday periods. The two midday periods, 10 a.m. to 2 p.m. and 2 p.m. to 6 p.m., serve as the reference groups.

For each city I compare the difference in morning search versus midday search before the fire to the corresponding difference in morning versus midday search after the fire. In other words, I test for a negative difference-in-differences in search in response to a negative difference-in-differences in price dispersion.

It is important to be clear about the identifying assumption underlying these comparisons, and what would potentially violate it. I am primarily interested in whether the difference between morning and midday search after the fire was less than the difference between morning and midday search before the fire:

$$S_{MORNING,AFTER} - S_{MIDDAY,AFTER} < S_{MORNING,BEFORE} - S_{MIDDAY,BEFORE} \quad (1)$$

where S represents search[11]. I am using a proxy for price dispersion based on time-of-day effects, in lieu of direct measurement, so the identifying assumption is:

$$PD_{MORNING,AFTER} - PD_{MIDDAY,AFTER} < PD_{MORNING,BEFORE} - PD_{MIDDAY,BEFORE} \quad (2)$$

where PD represents price dispersion. That is, I assume that the difference in price dispersion between the early morning and midday periods before the fire was larger than the difference in price dispersion between the early morning and midday periods after the fire, i.e. that price dispersion within the day became relatively more uniform after the fire. In other words, I am assuming that intraday price dispersion did not systematically fluctuate in a similarly synchronous but even more extreme way after the fire than had occurred before the fire. If the assumption is correct, I can attribute the difference-in-differences in search to an underlying difference-in-differences in price

[11] I make a similar comparison replacing the early morning (first intraday) period with the evening (fourth intraday) period instead, and comparing it to the midday periods.

dispersion.

The universal relationship between price cycles and price dispersion ensures that the right hand side of Equation 2 is large. Relenting phase price increases in Toronto were about 20 U.S. cents per gallon on average during the sample period, and as high as 30 U.S. cents per gallon, all staggered over a short few hours in the early morning periods. Meanwhile, undercutting phases in Toronto are known to produce price differences of generally only a few pennies across stations (Noel [2007b]). In contrast, the left hand side of Equation 2 is almost surely small, consistent with other gasoline markets in Canada and the United States that do not have daily retail price cycles. To the extent that prices change little within the day, the left hand side would be close to zero.

Equation 2 shows that only intraday price dispersion differences matter for the identifying assumption, price dispersion levels in their own right do not. I need not assume that price dispersion was zero after the fire or even that it was lower. Any common level of price dispersion or common level of consumer search that is intraday-invariant, i.e. that does not change from morning to midday to evening (e.g. individual consumer, firm, or city characteristics) is simply differenced out of the analysis by Equations 1 and 2. Similarly, any factor that is intraday-variant but present both before and after the fire (e.g. daily driving and demand patterns) is also differenced out of the analysis by Equations 1 and 2.

So what would violate the assumption? If midday price dispersion levels after the fire were similar to what they were before the fire, firms in Toronto or London post-fire would have to start randomizing their prices up or down by very large amounts each morning, large enough to create an even larger morning price gap between stations than had occurred when cycles were present. They would then have to retract those price changes a few hours later, each day, and resume competing within a price distribution as tight as it had been in the undercutting phases of the cycle. (Note that individual stations must randomize the direction of the price change each morning – not systematically increase or decrease them – or there would be a visible price cycle in the data which we do not observe.) If instead midday price dispersion levels after the fire were higher than before (the likely case given that price dispersion in an undercutting phase tends to be very low), it only means that firms would have to randomize prices by even greater amounts each morning to compensate for the higher midday price dispersion and keep the left hand side of

Equation 2 larger than the right. Such intraday price dispersion variation in the absence of price cycles, even putting aside the lack of an offerable explanation, has not been documented in gasoline markets without price cycles to my knowledge. Figures 3 and 4, which show a general stabilizing of average prices after the fire, also suggest that this is especially unlikely in the study markets.

Assuming the identifying assumption holds, interacting intraday periods with the before/after the fire variable is a valid proxy for the underlying changes in price dispersion. While direct measurement would be more ideal, the use of a proxy as a right hand side variable is not unusual in the price-dispersion-consumer-search literature, and is actually the norm. All the studies discussed above that examine the usual relationship of how consumer search affects price dispersion all rely on a proxy for consumer search on the right hand side and a direct measure of price dispersion on the left. In this study, I simply do the reverse – I use a proxy for price dispersion on the right hand side and a direct measure of search on the left.

Even still, the proxy that I use has an important advantage over many other proxies encountered in the literature. The nature of most proxies is that authors are limited to making only cross-sectional, single-difference comparisons between a high and low value of the proxy, e.g. when search costs are expected to be high versus when they are expected to low. Identification in these single-difference comparisons is relatively weaker and more subject to omitted variables biases. In contrast, the proxy I use here matches changes in price dispersion not along one but along two dimensions – changes in price dispersion across the periods within a day and also changes in price dispersion from before to after the fire – lending itself to a two-level difference-in-differences framework for improved identification.

III.2 A Direct Measure of Consumer Search

The second challenge in studying the reverse effect of price dispersion on consumer search is that a direct measure of search is difficult to find. Fortunately, a direct measure of search is available in the retail gasoline industry, using information from GasBuddy.com. GasBuddy is a popular gasoline-price-reporting website that allows consumers who perform on-the-street searches of gasoline prices to report the results of their search to a central website for the use and benefit of the greater public. I measure search as the number of gasoline price reports reported to the GasBuddy.com website by consumers who perform those primary on-the-street searches.

Search information from GasBuddy.com has been used in several previous studies – Lewis and Marvel [2011] use the number of pageviews to the national GasBuddy website from users of a particular browser toolbar while Byrne et al. [2015], like this paper, use the number of price reports submitted to the website. Both measures – pageviews received by the GasBuddy website and the number of price reports submitted to the GasBuddy website – are potential direct measures of consumer search. For my application, however, the latter measure has several important advantages.

First, unlike pageview data which is only available on a regional or national basis, a search measure based on price reports is available on a city specific basis. This matters because the impact on consumer search I am investigating should only occur in certain cities after the fire, and a city-specific measure is needed to isolate the treatment effects of interest. A measure of search based on regional or national pageviews would only wash out the effects of interest in aggregate averages.

Second, and unlike pageview data from after-the-fact online searches, my price-reporting measure is a truly contemporaneous measure of search with an exact timing match between the time a search is recorded and the time the prices observed in that search were actually in effect on the street.[12] The timing match is critical in my high-frequency application because the impact on consumer search I am investigating should be different for different hours of the day, and a contemporaneous and matched measure is needed to isolate the high-frequency treatment effects of interest. A measure of search based on pageviews after-the-fact contains a problematic variable lag of up to 24 hours from the time that prices (and price dispersion) were in effect on the street to the time that the search based on those prices is recorded, which would only mask the effects of interest.

The primary concern about a price-reporting measure of search is that, as it captures only a subset of search, it may not be representative of all consumers. Similar concerns have been voiced about pageviews, and this is fair. Since consumers that search and report the results of their searches to GasBuddy are likely to be more informed or more elastic than a typical consumer, my

[12] Reporters record both the prices they observe and the times at which they observed them. In fact, the price and search data are one and the same – price is calculated as the average price within an intraday period, and search is based on the number of price reports over which that average price is calculated.

results are likely most reflective of these types of consumers.

Having said this, there are several reasons why representativity should be of limited concern. First, and perhaps surprisingly, price reports to the popular GasBuddy site make up a substantial portion of all site usage. GasBuddy reports 49 million price reports per month and about 340 million pageviews per month on its main website. Since a price report itself requires one to two pageviews, the price reporting function makes up between one-sixth and one-third of all activity on the site. The sample size is larger that might be expected at first glance.

Second, there is evidence of a strong positive correlation between my measure of search and a proxy for all search. It is widely accepted that consumers search for gasoline prices more often when they are on the street and driving (or about to start driving) and, similarly, I find that my measure of consumer search is also highest at these times. Even in the presence of stable daily prices, the intraday correlation between the reporting-based measure of search and all consumer search (the latter proxied by hourly traffic counts) is 0.52.[13]

Third, the evidence shows that consumers who report prices to GasBuddy are reporting results from a broader price search and not just selectively reporting just 'low' or 'surprising' prices. Atkinson [2008] and Atkinson et al. [2014] compare GasBuddy price data to more comprehensive or systematic price datasets, in the context of cycling and non-cycling markets respectively, and show that the GasBuddy data does a good job representing the complete distribution of prices.[14]

Finally, GasBuddy is ultimately a social information exchange and the entire point of it is to help others learn what current gasoline prices are. Since the social value of reporting the results of one's search to the website is directly related to the number of other consumers currently searching for gasoline prices, it is not surprising that reporting-based search and all search are strongly correlated.

One final set of concerns relates to the interpretation of the measure itself. One concern is that since reporting is a technically separate function from searching, reporting may be influenced by non-search factors as well as search factors. This is possible, and allowable, as long as my reporting-

[13] Since I do not seek to place a cardinal interpretation on my estimates (i.e. that a 1% increase in price dispersion results in an X% increase in search), only a positive correlation between all search and my measure of search is needed.

[14] Atkinson [2008] and Lewis and Marvel [2011] also find that consumers who report prices tend to report many prices on a given day rather than just one or two, consistent with a broader search and not with selective reporting.

based measure of search remains positively correlated with all search, as I found above, and that the difference between them (the proxy error) is uncorrelated with price dispersion.[15] The fact that consumers who report prices to GasBuddy tend to report many prices on a given day rather than just one or two (Atkinson [2008] and Lewis and Marvel [2011]) also supports the interpretation that reporting activity largely reflects search motives. I take this interpretation going forward.

A final related concern is that, at least at first glance, a reporting based measure of search may not be the search-with-the-intention-to-purchase most often thought about by economists. My own view is that it is, once viewed in its proper social context. As mentioned, the purpose of GasBuddy is to disseminate price information to other consumers to assist them in their own searches, and the value of reporting price information to the site increases when there are more consumers searching on the site. In this light, a reporting-based measure is in fact a search-with-the-intention-to-purchase, where 'intention' is appropriately expanded to include not only one's own intention, but also the intention of the other users of the site that reporters wish to help.

III.3 Estimation

The main estimating equation is:

$$\begin{aligned} \ln(ReportCount)_{jst} &= \zeta_{j0} + \left(\sum_{s=2}^{4} \phi_{js}^{IDP}\right) + \zeta_{j1} AFTER\_REFINERY\_FIRE_t \\ &+ \left(\sum_{s=2}^{4} \zeta_{js} AFTER\_REFINERY\_FIRE_t * \phi_{js}^{IDP}\right) \\ &+ \sum_{d=2}^{7} \phi_{jd}^{DOW} + \sum_{m=2}^{3} \theta_{jm} + \eta_{jst} \end{aligned} \tag{3}$$

where $ln(ReportCount)$ is the natural log of the number of price reports in city j in intraday period s on day t, with $j = \{Toronto, London\}$ and $s = \{1..4\}$, and where $s = 1$ is the early morning period and $s = 4$ is the evening period. $AFTER\_REFINERY\_FIRE_t$ is an indicator variable equal to one after the fire, and ϕ_{js}^{IDP} is a set of fixed effects for each intraday period s. The first intraday period variable, ϕ_{j1}^{IDP} is the omitted variable. The ϕ_{jd}^{DOW} are day-of-the-week fixed effects with

[15] The only other obvious motive is that reporters earn points that can be used to enter a draw for a free tank of gas. Since points are awarded based on the number of price reports, and not on what prices actually are, this motive is by construction uncorrelated with price dispersion.

ϕ_{j1}^{DOW} (Sunday) being the omitted variable, and the θ_{jm} are monthly fixed effects with January 2007 being the omitted variable. I also present alternate specifications replacing the monthly fixed effects with a time trend T_t instead. To allow for a learning and adjustment period, I exclude the first 72 hours following the fire. The ζ_{j0} are city specific constants and the η_{jst} are independent and identically distributed error terms. Robust standard errors are calculated and presented.

I estimate city-specific difference-in-differences regressions as given in Equation 3, and also a pooled version of Equation 3 in which I combine both cities into a single regression according to their respective sizes. I test whether the ζ_{j1} and $\zeta_{j1} + \zeta_{js} * \phi_{js}^{IDP}$ are significantly different from zero (the change in search intensity in the first intraday period and other intraday periods after the fire), and whether each $\zeta_{js} * \phi_{js}^{IDP}$, $s \neq 1$, is significantly different from zero (the difference in the differences between early morning period search and other intraday period search before and after the fire). The null hypothesis is that the non-linear pattern of search intensity across the cycle (extremely high, very low, very low, low) did not change in any city after the fire. The alternative is that consumer search fell disproportionately in the first intraday period and potentially in the fourth intraday period as well, all relative to midday.

To confirm the results from the difference-in-differences analysis, I also present an alternative test of the effect of price dispersion on consumer search by examining how volatility in consumer search changed over the course of a day after the fire. The question is whether consumer search became more uniform, i.e. more consistent over the course of a day, when price dispersion also became more uniform in the absence of the cycles. To test this, I calculate the coefficient of variation in consumer search across the four intraday periods of a day in each city, before and after the fire, and compare them:

$$\begin{aligned} H_O \quad &: \quad \left(\frac{\sigma_{\ln ReportCount_{jst}}}{\mu_{\ln ReportCount_{jst}}} \mid AFTER_t = 1 \right) = \left(\frac{\sigma_{\ln ReportCount_{jst}}}{\mu_{\ln ReportCount_{jst}}} \mid AFTER_t = 0 \right) \\ H_A \quad &: \quad \left(\frac{\sigma_{\ln ReportCount_{jst}}}{\mu_{\ln ReportCount_{jst}}} \mid AFTER_t = 1 \right) < \left(\frac{\sigma_{\ln ReportCount_{jst}}}{\mu_{\ln ReportCount_{jst}}} \mid AFTER_t = 0 \right) \end{aligned} \tag{4}$$

where $\sigma_{\ln ReportCount_{jst}}$ and $\mu_{\ln ReportCount_{jst}}$ are the standard deviation and mean, respectively, of consumer search in city j in intraday period s and period t, and $AFTER_t$ is a synonym for

$AFTER\_REFINERY\_FIRE_t$.

Under the alternative hypothesis, consumer search across the four intraday periods becomes more uniform and the coefficient of variation falls. The mechanism is as follows. Under the alternative and in the presence of cycles, there are two main sources creating a non-uniformity in search over the course of the day. First, there is the natural daily pattern of consumer search, due to normal traffic flows, which results in higher search in the first intraday period (early mornings) and fourth intraday period (evening). The first intraday period includes the early morning rush hour and the fourth intraday period includes the latter part of the evening rush hour. Search in the fourth intraday period is also mechanically higher because it encompasses a six hour window instead of four. The second source of non-uniformity under the alternative comes from the universal pattern of price dispersion generated by the price cycles. The large relenting-phase spikes in price dispersion happen to coincide with the first (early morning) intraday periods, magnifying the extent of consumer search in those periods. Price dispersion can also increase in the fourth intraday period (evenings) magnifying the extent of consumer search then. Taken together, the natural daily search pattern and the cycle-induced price dispersion pattern compound to create a more U-shaped pattern of consumer search under the alternative and in the presence of cycles. Under the alternative and in the absence of price cycles, only the natural daily search pattern remains and search becomes relatively more uniform. The coefficient of variation should fall after the fire under the alternative.

In contrast, under the null hypothesis, price dispersion does not affect consumer search, only the natural daily search pattern is present with or without cycles, so there should be no change in the coefficient of variation after the fire.

IV Results

IV.1 Preliminaries

First, I simply establish that daily-period price cycles were present in Toronto and London before the fire, and then came to a stop immediately after the fire. This is readily apparent in Figures 1 and 3, above. Prior to the fire, the amplitude of the cycle was 5.5 cents per liter in Toronto and 5.2 cents per liter in London. The period in each case was one day, or four intraday periods. The asymmetry in each case, measured as the length of an undercutting phase divided by the length of a relenting phase, was equal to 4.0.

The presence and then absence of cycles can be tested statistically by comparing the median price change prior to the fire to the median price change after, using a technique suggested by Lewis [2009]. When prices rise by large amounts and fall by small amounts, as occurs with retail gasoline price cycles, the median price change tends to be relatively large and negative. Table II shows that the median price change before the fire was -0.73 and -1.05 cents per liter in Toronto and London respectively. After the fire, the median price change reverted to close to zero, 0.00 in Toronto and 0.02 cents per liter in London.

An alternate method examines changes in price volatility and is based on comparing average absolute price changes. The average absolute price change was 2.59 and 2.63 cents per liter in Toronto and London respectively before the fire, but only 0.47 and 0.33 cents per liter after, consistent with the loss of volatile intraday cycles.[16]

Average absolute price changes calculated on an intraday basis provide some additional support for the assumption that price dispersion was likely to be more uniform after the fire than before. The average absolute price change within the four intraday periods before the fire was 5.5, 0.3, 1.4, and 3.4 cents per liter in Toronto, respectively, and 5.0, 1.6, 1.6, and 2.3 cents per liter in London (with the first intraday period always showing increases and later intraday periods almost always showing decreases). After the fire, the average absolute price change fell to 0.9, 0.1, 0.1, and 0.1 cents per liter in Toronto and 0.6, 0.3, 0.5 and 0.4 cents per liter in London (with a relatively even mix of increases and decreases across intraday periods).[17] Noting that, in the post-fire world, the first estimate in each set represents a price change from the night before and the remaining three estimates refer to intraday price changes, the estimates show that intraday price changes in the post-fire world were small and, in Toronto especially, exceptionally close to zero. Like other markets without cycles, this likely reflects a situation where intraday price changes were small and few.[18]

[16] For the identifying assumption to be violated, a condition is that the unobserved average *station-specific* absolute price change would have to have increased at the same time that the city-average absolute price change fell so strongly.

[17] The slightly higher values at the start of a new day are expected as this is when most rack price changes take effect.

[18] The only other possibility is less likely – that there continued to be many large price changes (and perhaps

As an additional preliminary step, I present a set of regression discontinuity specifications to query whether consumer search changed overall after the fire. These are simple regressions that restrict $\zeta_{js} = 0 \; \forall \; s \neq 1$, and are only meant to provide context to the reader on overall search changes post-fire. By themselves they do not offer evidence of a relationship between price dispersion levels and consumer search, since the available data cannot determine whether overall price dispersion levels, averaged across the whole day, decreased or increased overall after the fire.

The results show that in each city there was a decline in consumer search overall occurring right around the time with the fire. Table III reports results in Specifications (1) through (6), with the $AFTER\_REFINERY\_FIRE$ coefficient being the coefficient of primary interest. (For readability, I replace ϕ_{j2}^{IDP} with $SECOND\_INTRADAY\_PERIOD$, ϕ_{j3}^{IDP} with $THIRD\_INTRADAY\_PERIOD$, and ϕ_{j4}^{IDP} with $FOURTH\_INTRADAY\_PERIOD$ in the tables.) Specification (1) shows that consumer search statistically significantly decreased in Toronto by 12.0% (coefficient -0.128) after the fire. Specification (2) shows that consumer search significantly fell in London as well, by 12.5% (coefficient -0.134). Specification (3) combines the two cities together in a pooled regression and shows a similar result. Specifications (4) through (6) are the corresponding specifications that use a time trend in place of monthly indicator variables and produce a similar result as well.[19] Though potentially consistent with a decrease in price dispersion overall, the regressions are not evidentiary and simply provide a baseline – an average decline in search of about 12% – against which we can compare the relative intraday consumer search changes that are of primary interest.

IV.2 Difference-in-Differences Analysis

The relevant question is whether that average reduction in search was disproportionately concentrated in the early mornings periods (and potentially in the evening periods as well), compared to midday. For this, I turn to the main results which are contained in Table IV, and consist of a series of difference-in-differences specifications. Specification (1) is for Toronto, Specification (2) is for London, and Specification (3) is the pooled regression.

even greater price dispersion volatility after the fire than before) but, for some reason or by chance, the city-average absolute price change suddenly became very stable.

[19] The point estimates are very similar, but the $AFTER\_REFINERY\_FIRE$ coefficient is not statistically significant. Alternate specifications that include both monthly fixed effects and time trends produce nearly identical results and regain statistical significance on all coefficients.

In all specifications, the $AFTER\_REFINERY\_FIRE$ coefficient shows a disproportionately large change in search in the early mornings after the fire, as expected. In Specification (1) for Toronto, I find a coefficient of -0.22, implying a 20.0% reduction in consumer search in the early morning hours after the fire ($exp(-0.22) - 1 = 20.0\%$). The coefficients on the midday periods, $SECOND\_INTRADAY\_AFTER$ and $THIRD\_INTRADAY\_AFTER$ are positive (in one case statistically significant), meaning that midday search did not fall by as much as in the early morning periods after the fire. Summing the relevant coefficients (of each interaction term with the $AFTER\_REFINERY\_FIRE$), I calculate that consumer search fell by only 6.9% ($exp(-0.22 + 0.15) - 1$) in the second intraday period and 2.0% ($exp(-0.22 + 0.20) - 1$) in the third intraday period, both insignificantly different from zero. Search fell significantly after the fire in the fourth intraday period as well, by 19.7% ($exp(-0.22 + 0.03) - 1$).

Specification (2) shows a similar result for London. I find a statistically significant reduction in consumer search in the early morning hours after the fire compared to before the fire of 25.0% (coefficient on the $AFTER\_REFINERY\_FIRE$ indicator equal to -0.287). The midday coefficients $SECOND\_INTRADAY\_AFTER$ and $THIRD\_INTRADAY\_AFTER$ are positive (in one case statistically significant), showing that midday search did not fall by as much as in the early morning periods. Summing the relevant coefficients, I find that search in the second period after the fire was actually higher by an insignificant amount and search in the third period was lower. The coefficient on $FOURTH\_INTRADAY\_PERIOD$ is positive showing that consumer search fell in the fourth intraday period but by a lesser amount than in the first intraday period. Summing the relevant coefficients, I find that search intensity fell by 11.1% in the fourth (evening) intraday period.

Specification (3) pools the two cities together, with similar results. There was a large and significant reduction in search intensity of 20.5% focused on the early morning periods after the fire ($exp(-0.229) - 1 = 20.5\%$). There was an insignificantly small change in search intensity in the two midday periods ($exp(-0.229 + 0.189) - 1 = 3.9\%$ and $exp(-0.229 + 0.171) - 1 = 5.6\%$), and a significant but intermediate change in search intensity in the evening period ($exp(-0.229 + 0.044) - 1 = 16.8\%$). The difference in the change between early morning search and early midday search before to the fire is significant at the 10% level.[20]

Specifications (4) through (6) repeat Specifications (1) through (3) using a time trend instead of monthly indicator variables. The results are nearly identical and all the main conclusions hold.[21] I conclude that when the large early morning spikes to price dispersion disappeared with the price cycles, the large early morning spikes in consumer search disappeared at exactly those same times as well.

As a robustness check, I also estimate a set of a linear-linear specifications instead of log-linear ones (not shown) and all results continue to hold. The increase in consumer search in the early morning period is even more pronounced in these specifications since search is highest in absolute terms during the early morning period, making a given percentage increase in the early morning period even higher when expressed in levels.

For context, Figure 5 shows the raw change in consumer search immediately before and after the fire in the two cities. The most comprehensive data is for Toronto (1404 price reports per day) and shows a clear reduction in the daily peaks of consumer search activity after the fire. The peaks correspond to the early morning period, the troughs correspond to the midday periods, and consumer search in the evening period is intermediate between the two (generally between one-half and two-thirds up the upward portion between trough and peak). It took only a few days for consumers to recognize the change in the price dispersion pattern and downward adjust their search activity. The data for London, while thinner, shows a similar overall compression of consumer search after the fire.

IV.3 Uniformity Analysis

An alternate test of the effect of price dispersion on consumer search is based on comparing the degree of uniformity in consumer search after the fire compared to before the fire. Since the natural daily search pattern was magnified by the cycle-induced price dispersion pattern before the fire, but only the natural search pattern remains after, the coefficient of variation should fall under the alternative hypothesis.

[20] The difference in change between early morning search and late midday search is significant at the 11% level.

[21] Including monthly fixed effects and time trends in the same specifications produce nearly identical results.

Table V presents results and confirms my main conclusion. The coefficient of variation for Toronto fell from 0.058 before the fire to 0.047 after, a statistically significant decrease at the 2% level. The coefficient of variation for London also fell substantially, from 0.063 before the fire to 0.050 after, significant at the 2% level. I conclude that intraday consumer search did indeed become relatively more uniform (i.e. more consistent over the day) just as price dispersion became relatively more uniform over the day in the absence of the large early morning cycle-induced price dispersion spikes.[22]

To see the increase in uniformity in its raw form, I report the pattern of search given by the average value of *ReportCount* across the four intraday periods. In Toronto, the pattern was 331, 188, 168, and 305, before the fire. After the fire, the pattern was 265, 175, 165, and 250. In London, the search pattern was 49, 28, 34, and 37 prior the fire, and 36, 32, 27, and 33 after. In each case, search became more uniform after the fire than before, with the largest reduction occurring in the early morning intraday period.

IV.4 Falsification Exercise

For an additional layer of evidence, I add a third city to the sample to use as a falsification exercise – the city of Vancouver, British Columbia. Vancouver is the only city in the available data known to have daily cycles but be far removed from the fire. Vancouver is on the west coast of Canada, approximately 2,800 miles away from Nanticoke, and is served by Western Canadian refineries and not by Nanticoke. There was no sudden stop to the price cycles or the cycle-induced daily pattern of price dispersion in this city, and there should be no disproportionate change in consumer search in the early morning periods.

Figure 6 confirms that the cycles were still present in Vancouver before and after the February 15th date of the fire, and that retail prices were similar in nature as well. Margins were also similar before and after. The median price change in Vancouver was -0.18 cents per liter prior to the February 15th fire, and -0.45 cents per liter after. The average absolute price change was 1.61 cents per liter before and 1.38 cents per liter after. With no change in the cycles, I expect no sudden and disproportionate shock to early morning consumer search after February 15th, as there

[22] Comparing standard deviations instead of coefficients of variation produces similar and significant results.

was in Toronto and London.

Table VI shows the regression results for Vancouver – the preliminary regression discontinuity results in Specifications (1) and (2) and the main difference-in-differences results in Specifications (3) and (4).

The regression discontinuity specification in Specification (1), which includes monthly indicator variables but not a time trend, shows a statistically insignificant and general decline in consumer search in Vancouver after February 15th compared to before. Perhaps surprisingly, the magnitude of the $AFTER\_REFINERY\_FIRE$ coefficient appears not too much lower that of Toronto and London in Table III. However, upon closer inspection, it is clear that the Vancouver change did not occur suddenly after the fire but is attributable to a longer term underlying trend of declining search that was not present in Toronto and London. This can be seen in Specification (2) which includes a time trend instead of monthly indicator variables and shows that the point estimate on the $AFTER\_REFINERY\_FIRE$ coefficient for Vancouver largely vanishes – it now falls by two-thirds and is highly statistically insignificant.[23] In contrast, when the time trend was included in Toronto and London, the magnitude of the effects in those cities remained unchanged. In other words, the reduction in search in Toronto and London was sudden with the fire, whereas the gradual decline in consumer search in Vancouver was not.

Again, the main question is whether search fell disproportionately in the early morning hours and potentially in the evening hours in Vancouver after February 15th, as occurred in Toronto and London. Specifications (3) and (4) test this, without and with the time trend respectively. I expect there to be no disproportional shock to early morning search since the cycles in Vancouver were unaffected and price dispersion patterns did not change. The results confirm this – in both cases, there is no statistically significant disproportional decrease in consumer search in the early morning periods. In Specification (4) which includes the time trend, the point estimate on $AFTER\_REFINERY\_FIRE$ is just -0.02, corresponding to a decrease in consumer search of just 2%, and indistinguishable from zero. Contrast this to early morning declines in consumer search of 20-25% for Toronto and London. The midday and evening coefficients for Vancouver are

[23] An alternate specification that includes both monthly dummies and a time trend produces a similar result to Specification (2).

all insignificantly different from zero and insignificantly different from each other, showing that the decline in consumer search in Vancouver was equally spread across all four intraday periods, consistent with a longer term declining trend.[24] This is again in contrast to Toronto and London which showed a disproportional decrease in early morning search immediately after the fire.

It is also possible to test for differential early morning effects across the three cities by combining Toronto, London, and Vancouver into a single system, essentially a three-level difference-in-differences-in-differences framework (the first difference being the before/after, the second difference being the intraday period, and the third difference being the affected/unaffected city). In cross-equation coefficient tests, I find that the post-February-15th change in the difference between early morning and midday search in Toronto and London on one hand, and the change in the difference between early morning and midday search in Vancouver on the other, were significantly different from each other at the 10% level. That is, there is a statistically significant difference in the differences in the differences (p-value = 0.098 for Toronto, 0.059 for London). The results confirm that, while the gradual decrease in search that occurred in Vancouver came proportionally from all periods within the day, the sudden decrease in consumer search in Toronto and London came right after the fire and was focused on the early morning intraday periods, exactly when the largest shocks to price dispersion in those cities occurred.

As an additional check, I perform the uniformity test again, this time applied to Vancouver. Under the null, there should be no meaningful change in intraday uniformity and this is exactly what I find. The coefficient of variation for Vancouver was unchanged, from 0.074 before the fire to 0.070 after, a statistically insignificant difference with a p-value of 0.77.[25] When I combine all three cities into a single system and test for the difference-in-differences in the coefficients of variation, comparing the difference in the change in the coefficient of variation in Vancouver on one hand and the change in the coefficients of variation in Toronto and London on the other, I find a statistically significant difference in the differences at the 10% level. I conclude that the distribution of search activity in Toronto and London across the day did indeed become more uniform after the cycles ceased, but that there was no such redistribution of intraday search activity in Vancouver,

[24] Results are similar and conclusions do not change when including both monthly indicator variables and a time trend in the same specification.

[25] Comparing standard deviations instead of coefficients of variation produces similar and insignificant results.

as expected if the shock to price dispersion caused the shock to consumer search in Toronto and London.

IV.5 Alternate Hypotheses

The natural experiment, two testing frameworks, and a falsification exercise all point to the same conclusion – that the change in price dispersion impacted consumer search in ways predicted by the theory.

One potential concern about the interpretation of the results is that the cessation of cycles may not only have impacted price dispersion but also general price levels, and that a change in price levels might affect consumer search as well. While the elimination of the large early morning price dispersion spikes is the most obvious effect of the cessation of cycles, the fire did cause a small general price increase (Noel [2015]). If so, consumer search might be responding to overall price levels instead of price dispersion. Fortunately, the concern is easily handled by the difference-in-differences framework. I am essentially comparing search in the early morning intraday period to search in the later intraday periods, before and then after the fire. By definition, a general increase in price levels is intraday-invariant, so any change in consumer search that might be caused by a change in price levels is simply differenced out in the analysis. Putting aside the fact that price changes after the fire were not historically unusual, and positive (which would presumably lead to more search and not less search as I found), only differential intraday changes can potentially matter, so this concern is easily set aside.

Another concern might be that the temporary supply shortage caused some general change in consumer search behavior. Although no general change in consumer behavior was apparent, any such change would presumably affect search all through the day, and not just in the early morning hours, meaning that this effect would also be intraday-invariant and differenced out of the analysis.

One last concern is that, instead of price levels, consumer search may be responding to short run changes in the price instead. Though unlikely ex ante, if consumer search is, for whatever reason, responding to very short term intraday price changes in their own right, rather than the changes in price dispersion caused by them, then the effects found here may be a combination of price dispersion effects and short run price change effects. After all, the large price dispersion spikes occur at the exactly the same time that most price increases also occur, meaning there is an intraday-variant change in the pattern of price changes at the same time as an intraday-variant change in the pattern of price dispersion.

The evidence, presented in Specifications (1) through (3) of Table VII, alleviates this concern by showing that consumer search predominantly responds to changes in price dispersion rather than changes in price. The specifications are based on the estimating equation:

$$\begin{aligned}\ln(ReportCount)_{jst} &= \theta + \lambda^{+}\Delta p^{+}_{j,t+1} + \lambda^{-}\Delta p^{-}_{j,t+1} + \sum_{i=0}^{N-1}\kappa_i^{+}\Delta p^{+}_{j,t-i} + \sum_{i=0}^{N-1}\kappa_i^{-}\Delta p^{-}_{j,t-i} \\ &\quad + \left(\sum_{s=2}^{4}\phi_{js}^{IDP}\right) + \sum_{d=2}^{7}\phi_{jd}^{DOW} + \sum_{m=2}^{3}\theta_{jm} + \sum_{j=2}^{3}\phi_j^{CITY} + \eta_{jst} \qquad (5)\end{aligned}$$

where $\Delta p^{+} = \max(0, \Delta p)$ and $\Delta p^{-} = \min(0, -\Delta p)$ represent positive and negative price changes respectively, and the lag length N is set at twelve to capture short run search dynamics over multiple cycles. I freely estimate κ_0^{+} and κ_0^{-} and set $\kappa_{1-3}^{+} = \kappa_1^{+} = ... = \kappa_3^{+}$ (the rest of the first day) and $\kappa_{4-11}^{+} = \kappa_4^{+} = ... = \kappa_{11}^{+}$ (the second and third days) for enhanced readability. I define κ_{1-3}^{-} and κ_{4-11}^{-} similarly, and my results are not meaningfully affected using other constrained ranges. Other variables are as previously defined.

Specification (1) is a constrained version of the regression that restricts each ϕ_{js}^{IDP} to be zero. In other words, it includes all lead, lagged, and current price changes but excludes all intraday indicator variables, recognizing that there are price changes, but essentially ignoring the regular time-of-day price dispersion pattern embodied in the price cycles. Specification (2) is also a constrained version of the regression, but this time restricting all $\lambda's$ and $\kappa's$ to be zero. That is, it includes all intraday indicator variables but excludes all lead, lagged, and current price changes, recognizing the cycle pattern but essentially ignoring by how much prices are changing. Specification (3) is the unconstrained regression that contains all the variables given in Equation 3.

The result in Specification (1) shows that – when ignoring the time of day – consumer search increases significantly when there is a contemporaneous price increase ($\Delta p_t^{+} = 0.078$), but does not respond to any other price change, positive or negative, lag or lead, at any other times. In other words, search does not react to price changes at all – except when prices are contemporaneously rising.

Of course, the vast majority of price increases occur in the early morning periods of the day as part of the relenting phase of the cycle, the same time that spikes to price dispersion occur, which leads to the simultaneity concern. This can be seen in Specification (2), which excludes lagged, lead, and current price changes and includes only the intraday variables (again using the more informative names in place of the ϕ_{js}^{IDP} in the table). Specification (2) shows that consumer search is significantly higher in the early morning periods (the omitted intraday variable) when price dispersion spikes and less at other times, echoing earlier results.

Specification (3) now puts both sets of variables together in the same regression. It is essentially a strength-of-effects comparison, and shows that consumer search at this high frequency is in fact being driven by the large price dispersion spikes that occur in the early morning intraday periods and not by price increases in their own right. I find that the coefficient on contemporaneous price increases now largely disappears (falling by over 70%) and is now statistically insignificant. In contrast, the intraday search effect remains significant and strong. I conclude that any intraday price increase that does not occur in the early morning period and is not associated with the large early morning price dispersion spike does not result in higher search. Only intraday price increases that are associated with the large early morning price dispersion spikes do. The result is intuitive – higher prices in their own right are unlikely to affect the benefits of consumer search, or consumer search itself, unless at the same time there are still lower priced stations around to buy from, i.e. consumer search is only useful when price dispersion is present.

V Conclusion

In this article, I examine the effect of a shock to equilibrium price dispersion on consumer search in retail gasoline markets. The question of how price dispersion affects consumer search, in contrast to the question of how consumer search affects price dispersion, has received little attention in the literature. The reason is likely due to two practical difficulties – the difficulty in finding a direct measure of search and the difficulty in finding a plausibly exogenous shock to price dispersion. I overcome both issues here, first using a direct measure of search based on primary on-the-street searches reported to a gasoline price reporting website, and second by exploiting a unique natural experiment in which a refinery fire caused decades-old retail gasoline price cycles to cease in several cities. Using the non-linear shock to price dispersion to identify effects on consumer search, I find that search decreased disproportionately at exactly those times of the day when price dispersion had disproportionately fallen after the cycles ceased, consistent with the theory. In other words, I find that consumer search responds positively to shocks to price dispersion.

I offer several contributions. First, this study adds to our understanding of the two-way relationship between price dispersion and search. Search activity not only affects price dispersion, as has been well documented, but price dispersion also impacts search, an issue rarely investigated. Second, by empirically confirming the reverse relationship and thus the theoretical result that consumer search and price dispersion are jointly determined, this study highlights a potential bias in previous studies that examine how consumer search affects price dispersion (the usual studied relationship) but do not account for the endogeneity. Understanding this two-way relationship can help us quantify and correct the potential bias contained in those other studies.

To my knowledge, this is the first study that examines the reverse relationship in a natural experiment setting. Exogenous shocks to price dispersion are difficult to find, but the question of how price dispersion affects search is no less interesting, and additional research is needed going forward. One direction for future research is to simply extend the analysis to other gasoline markets and other industries, to quantify effects, and to study any differences that may arise across industries and over time. Such research can help us better understand why consumers search, and what factors are most important for equilibrium levels of consumer search in different environments.

VI References

Allen, A, Alleweldt, F., Beteille, R., Blondin, D., Dewenter, R., Eckert, A., Fotis, P., Friedrich, A., Graham, R., Kara, S., Kimova, T., McSpedden-Brown, N., Noel, M., Perfors, H., Riviere, O., Schwarzer, S., West, D., 2014, 'Consumer Market Study on the Functioning of the Market for Vehicle Fuels from a Consumer Perspective'. *Report prepared for the Executive Agency for Health and Consumers acting on behalf of the European Commission*, pp. 1-398.

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

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

Barron, J.M., Taylor, B.A. and Umbeck, J.R., 2004, ‘Number of Sellers, Average Prices, and Price Dispersion’, *International Journal of Industrial Organization* 22(8), pp. 1041-1066.

Baye, M.R., Morgan, J., and Scholten, P., 2003, ‘The Value of Information in an Online Consumer Electronics Market’, *Journal of Public Policy and Marketing* 22(1), pp. 17-25.

Bloch, H. and Wills-Johnson, N., 2010, ‘Gasoline Price Cycle Drivers’, Working Paper, Curtin University of Technology.

Borenstein, S., Cameron, A., and Gilbert, R., 1997, ‘Do gasoline markets respond asymmetrically to crude oil price changes?’, *Quarterly Journal of Economics* 112, pp. 305-339.

Brynjofsson, E. and Smith, M.D., 2000, ‘Frictionless Commerce? A Comparison of Internet and Conventional Retailers’, *Management Science* 46(4), pp. 563-585.

Burdett, K. and Judd, K.L., 1983, ‘Equilibrium Price Dispersion’, *Econometrica* 51(4), pp. 955-969.

Byrne, D.P., 2015, ‘Gasoline Pricing in the Country and the City’, Working Paper, University of Melbourne.

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

Byrne, D.P. and de Roos, N., 2015, ‘Consumer Search in Retail Gasoline Markets’, *Journal of Industrial Economics*, pp. 183-193.

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

Chandra, A. and Tappata, M., 2011, ‘Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets’, *RAND Journal of Economics*, 42(4), pp. 681-704.

Dahlby, B. and West, D.S., 1986, ‘Price Dispersion in an Automobile Insurance Market’, *Journal of Political Economy* 94(2), pp. 418-438.

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

Diamond, P., 1971, ‘A Model of Price Adjustment’, *Journal of Economic Theory* 3(2), pp. 156-168.

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

Dubois, P. and Perrone, H., 2015, ‘Price Dispersion and Informational Frictions: Evidence from Supermarket Purchases’, Working Paper, Toulouse School of Economics.

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

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

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

Hong, H., and Shum, M., 2006, ‘Using Price Distributions to Estimate Search Costs’, *RAND Journal of Economics*, 37(2), pp. 257-275.

Hosken, D.S., McMillan, R.S. and Taylor, C.T., 2008, ‘Retail Gasoline Pricing: What do we Know?’, *International Journal of Industrial Organization*, 26(6), pp. 1425-1436.

Lewis, M. and Marvel, H.P., 2011, ‘When Do Consumers Search?’, *Journal of Industrial Economics*, 59(3), pp. 457-483.

Lewis, M., 2008, ‘Price Dispersion and Competition with Differentiated Sellers’, *Journal of Industrial Economics*, 56(3), pp. 654-678.

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), pp. 581-605.

Lewis, M., 2011, 'Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market', *Journal of Economics and Management Strategy*, 20(2), pp. 409-449

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

Lewis, M. and Noel, M., 2011, 'The speed of gasoline price response in markets with and without Edgeworth Cycles', *Review of Economics and Statistics* 93(2), pp. 672-682.

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

Milyo, J. and Waldfogel, J., 1999, 'The Effect of Price Advertising on Prices: Evidence in the Wake of 44 Liquormart', *American Economic Review* 89(5), pp. 1081-1096.

Nishida, M., and Remer, M., forthcoming, 'The Determinants and Consequences of Search Cost Heterogeneity: Evidence from Local Gasoline Markets', *Journal of Marketing Research.*

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), pp. 324-334.

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

Noel, M.D., 2008, 'Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria', *Journal of Economics and Management Strategy*, 17(2), pp. 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), pp. 582-595.

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

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

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

Orlov, E., 2011, 'How Does the Internet Influence Price Dispersion? Evidence from the Airline Industry', *Journal of Industrial Economics* 59(1), pp. 21-37.

Pennerstorfer, D., Schmidt-Dengler, P., Schutz, N., Weiss, C. and Yontcheva, B., 2015, 'Information and Price Dispersion: Evidence from Retail Gasoline', Discussion Paper DP10771, Centre for Economic Policy Research (CEPR).

Sherman, J. and Weiss, A., 2017, 'On Fruitful and Futile Tests of the Relationship between Search and Price Dispersion', *Economic Inquiry* 55(4), pp. 1898-1918.

Siekmann, 2017, 'Characteristics, Causes, and Price Effects: Empirical Evidence of Intraday Edgeworth Cycles', Discussion Paper 252, Dusseldorf Institute for Competition Economics (DICE).

Sorensen, A., 2000, 'Equilibrium Price Dispersion in Retail Markets for Prescription Drugs', *Journal of Political Economy* 108(4), pp. 833-850.

Stahl, D. O., 1989, 'Oligopolistic Pricing with Sequential Search', *American Economic Review* 79(4), pp. 700-712.

Stigler, G.J., 1961, 'The Economics of Information', *Journal of Political Economy* 69(3), pp. 213-225.

Varian, H. R., 1980, 'A Model of Sales', *American Economic Review* 70(4), pp. 651-659.

Valadkini, A., 2013, 'Seasonal Patterns of Daily Prices of Unleaded Petrol Across Australia', *Energy Policy* 56, pp. 720-731.

Walsh, P.P. and Whelan, C., 1999, 'Modelling Price Dispersion as an Outcome of Competition in the Irish Grocery Market', *Journal of Industrial Economics* 47(3), pp. 325-343.

Wang, Z., 2009a, 'Station Level Gasoline Demand in an Australian Market with Regular Price Cycles', *Australian Journal of Agricultural and Resource Economics*, 53(4), pp. 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), pp. 987-1030.

Zhao, Y., 2006, 'Price Dispersion in the Grocery Market', *Journal of Business* 79(3), pp. 1175-1192.

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

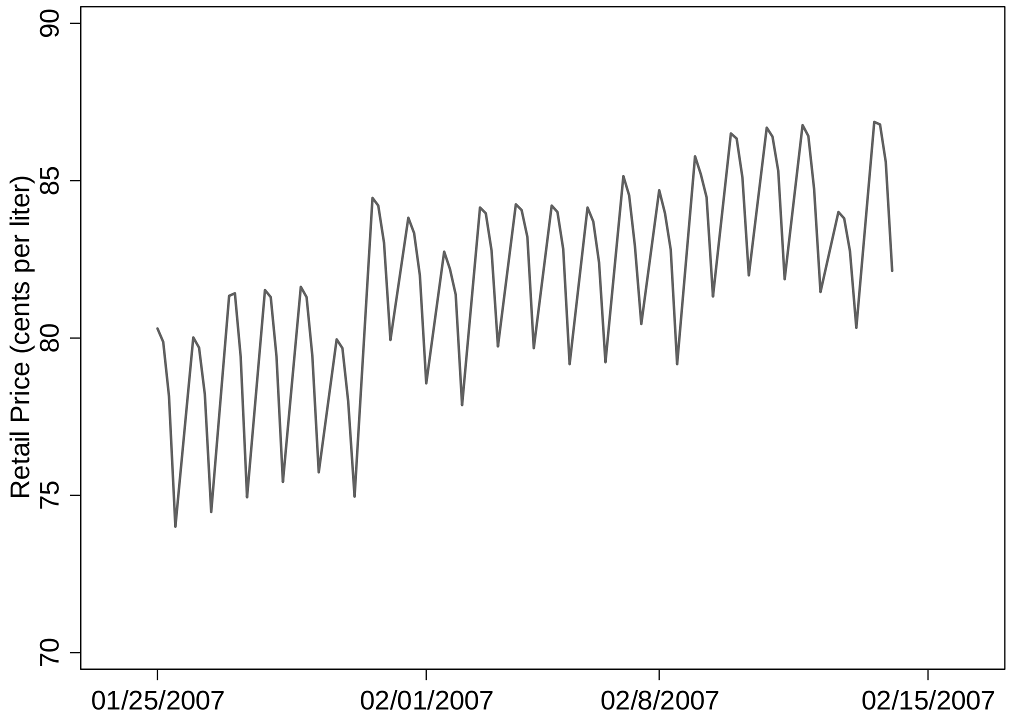

Figure 1
Daily Price Cycles in Toronto Before the Fire

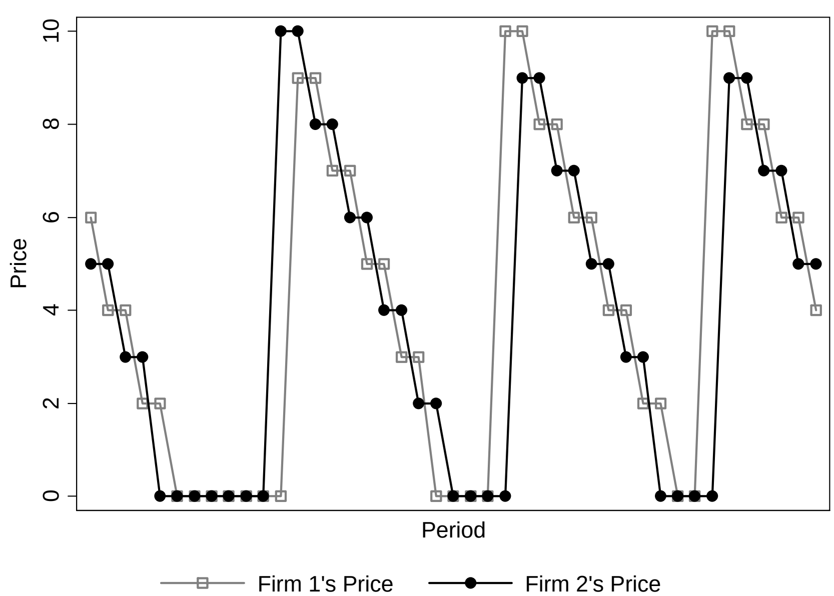

Figure 2
Price Dispersion Along a Theoretical Edgeworth Price Cycle

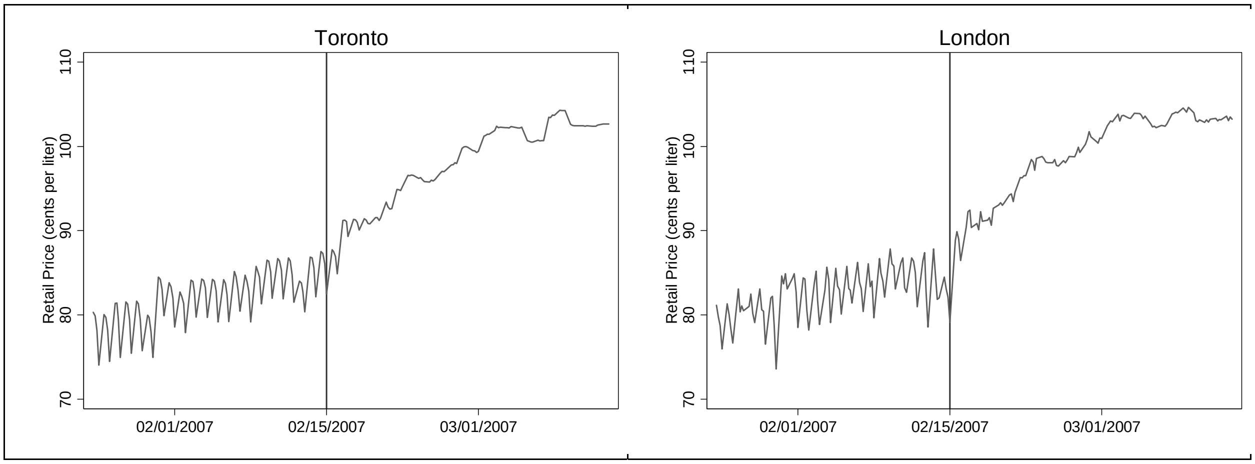

Figure 3
Retail Prices in Toronto and London Before and After the Fire

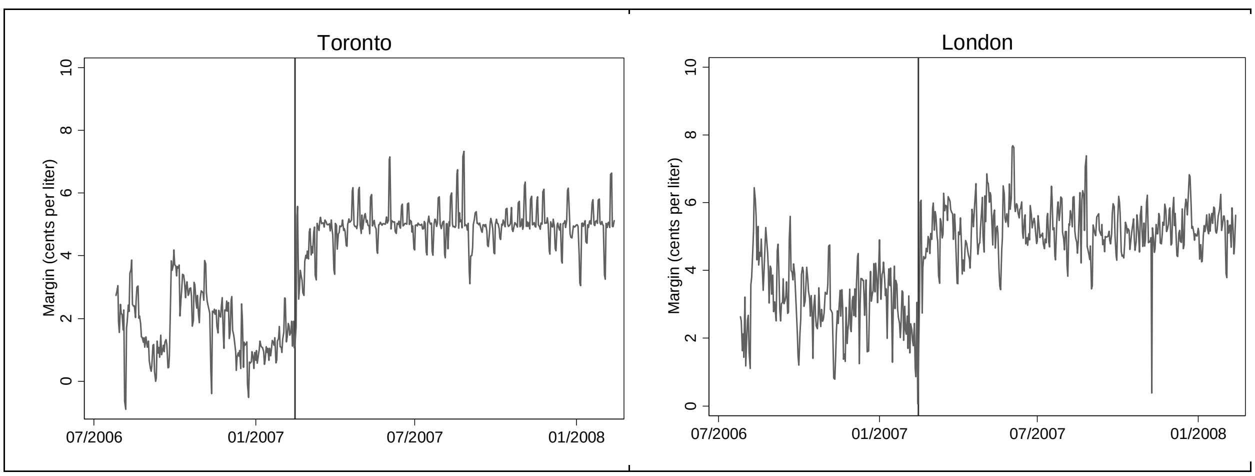

Figure 4
Retail Margins in Toronto and London Before and After the Fire

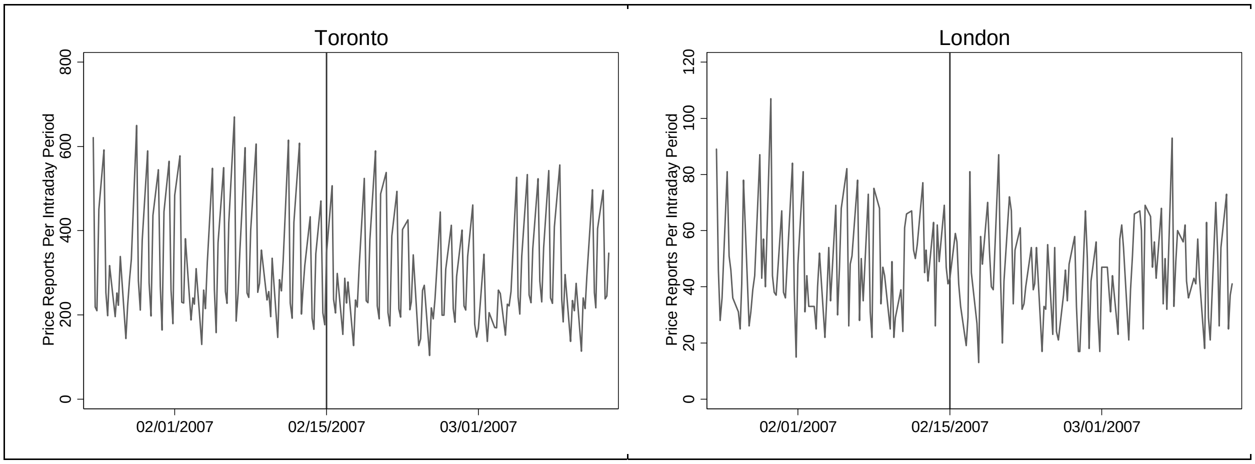

Figure 5
Intraday Consumer Search in Toronto and London Before and After the Fire

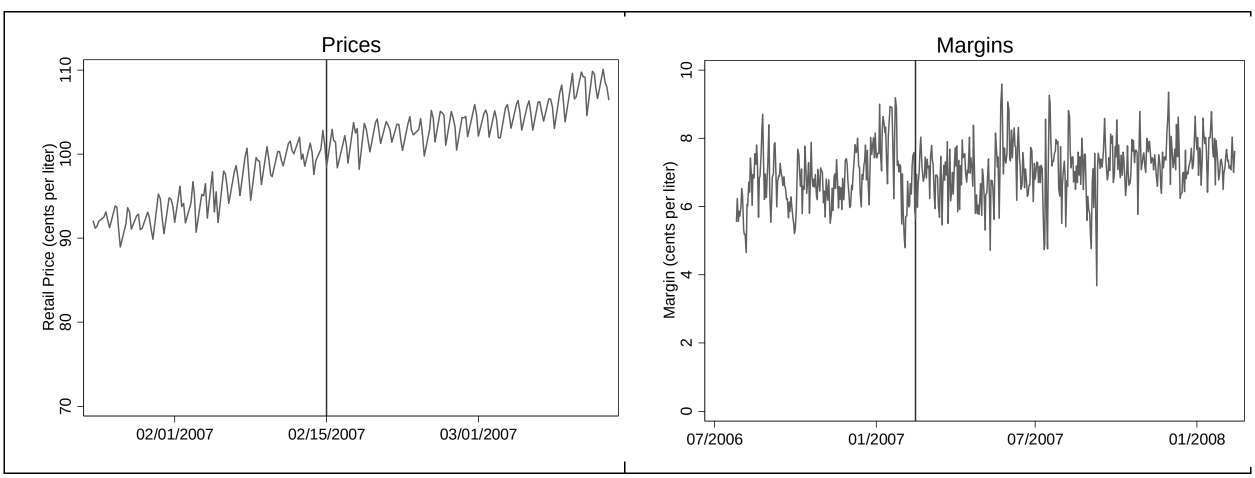

Figure 6
Retail Prices and Margins in Vancouver Before and After the Fire

TABLE I

SUMMARY STATISTICS

	Mean	Std. Dev.	Minimum	Maximum
Four-times-a-day Sample:				
PRICE (cents per liter)	92.48	9.99	71.83	110.09
REPORTCOUNT	131.56	150.09	2.00	670.00
Daily Sample:				
PRICE (cents per liter)	99.90	10.01	75.31	127.96
RACK PRICE (cents per liter)	62.76	7.40	45.30	83.70
TAXES (cents per liter)	32.40	3.13	28.96	38.35
REPORTCOUNT	584.37	619.55	83.00	2297.00

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

TABLE II

PRICE CHANGES BEFORE AND AFTER THE FEBRUARY 15TH FIRE

	Median Price Change		Average Absolute Price Change	
	Before	After	Before	After
Toronto	-0.73	0.00	2.59	0.47
London	-1.05	0.02	2.63	0.33

In Canadian cents per liter.

TABLE III

REGRESSION DISCONTINUITY RESULTS

Dep. Var.: ln(ReportCount)	Toronto (1)	London (2)	Combined (3)	Toronto (4)	London (5)	Combined (6)
SECOND_INTRADAY_PERIOD	-0.500**	-0.383**	-0.480**	-0.500**	-0.383**	-0.479**
	(0.056)	(0.071)	(0.054)	(0.056)	(0.072)	(0.054)
THIRD_INTRADAY_PERIOD	-0.588**	-0.339**	-0.547**	-0.588**	-0.339**	-0.547**
	(0.056)	(0.071)	(0.054)	(0.056)	(0.072)	(0.054)
FOURTH_INTRADAY_PERIOD	-0.070	-0.210**	-0.084	-0.070	-0.211**	-0.083
	(0.056)	(0.071)	(0.054)	(0.056)	(0.072)	(0.054)
AFTER_REFINERY_FIRE	-0.128**	-0.134*	-0.127**	-0.112	-0.110	-0.108
	(0.060)	(0.076)	(0.057)	(0.084)	(0.108)	(0.081)
Day-of-the-Week Indicators	Y	Y	Y	Y	Y	Y
Month Indicators	Y	Y	Y	N	N	N
Time Trend	N	N	N	Y	Y	Y
Num. Obs.	220	220	220	220	220	220
Adj. R-squared	0.513	0.199	0.501	0.514	0.189	0.502

Standard errors in parentheses. ** Significant at the 5% level. * Significant at the 10% level. The omitted intraday period is the first intraday period, from 6 a.m. to 10 a.m. The second intraday period is 10 a.m. to 2 p.m., the third intraday period is 2 p.m. to 6 p.m., and the fourth intraday period is 6 p.m. to 12 a.m. midnight.

TABLE IV
DIFFERENCE-IN-DIFFERENCES RESULTS

Dep. Var.: ln(ReportCount)	Toronto (1)	London (2)	Combined (3)	Toronto (4)	London (5)	Combined (6)
SECOND_INTRADAY_PERIOD	-0.566**	-0.566**	-0.562**	-0.566**	-0.566**	-0.562**
	(0.074)	(0.093)	(0.071)	(0.074)	(0.094)	(0.071)
THIRD_INTRADAY_PERIOD	-0.676**	-0.349**	-0.622**	-0.676**	-0.350**	-0.622**
	(0.074)	(0.093)	(0.071)	(0.074)	(0.094)	(0.071)
FOURTH_INTRADAY_PERIOD	-0.081	-0.284**	-0.103	-0.081	-0.285**	-0.102
	(0.074)	(0.093)	(0.071)	(0.074)	(0.094)	(0.071)
AFTER_REFINERY_FIRE	-0.222**	-0.287**	-0.229**	-0.206*	-0.264*	-0.209**
	(0.091)	(0.114)	(0.087)	(0.109)	(0.137)	(0.104)
SECOND_INTRADAY * AFTER	0.150	0.421**	0.189*	0.150	0.421**	0.189*
	(0.113)	(0.141)	(0.108)	(0.112)	(0.142)	(0.108)
THIRD_INTRADAY * AFTER	0.202*	0.024	0.171	0.202*	0.024	0.171
	(0.113)	(0.141)	(0.108)	(0.112)	(0.142)	(0.108)
FOURTH_INTRADAY * AFTER	0.026	0.169	0.044	0.026	0.169	0.044
	(0.113)	(0.141)	(0.108)	(0.112)	(0.142)	(0.108)
Day-of-the-Week Indicators	Y	Y	Y	Y	Y	Y
Month Indicators	Y	Y	Y	N	N	N
Time Trend	N	N	N	Y	Y	Y
Num. Obs.	220	220	220	220	220	220
Adj. R-squared	0.516	0.230	0.505	0.518	0.220	0.505

Standard errors in parentheses. ** Significant at the 5% level. * Significant at the 10% level. The omitted intraday period is the first intraday period, from 6 a.m. to 10 a.m. The second intraday period is 10 a.m. to 2 p.m., the third intraday period is 2 p.m. to 6 p.m., and the fourth intraday period is 6 p.m. to 12 a.m. midnight.

TABLE V
COEFFICIENT OF VARIATION IN SEARCH BEFORE AND AFTER FEB. 15TH FIRE

	Before	After	Difference	P-Value
Toronto	0.058	0.047	0.011	0.0176
London	0.063	0.050	0.013	0.0104

TABLE VI
FALSIFICATION EXERCISE - VANCOUVER

Dep. Var.: *ln(ReportCount)*	Regression Discontinuity (1)	(2)	Difference-in-Differences (3)	(4)
SECOND_INTRADAY_PERIOD	-0.582**	-0.581**	-0.592**	-0.591**
	(0.086)	(0.086)	(0.115)	(0.115)
THIRD_INTRADAY_PERIOD	-0.067	-0.066	-0.031	-0.030
	(0.086)	(0.086)	(0.115)	(0.115)
FOURTH_INTRADAY_PERIOD	-0.079	-0.077	-0.060	-0.058
	(0.086)	(0.086)	(0.115)	(0.115)
AFTER_REFINERY_FIRE	-0.119	-0.046	-0.094	-0.020
	(0.091)	(0.129)	(0.140)	(0.168)
SECOND_INTRADAY * AFTER			0.023	0.023
			(0.174)	(0.174)
THIRD_INTRADAY * AFTER			-0.081	-0.081
			(0.174)	(0.174)
FOURTH_INTRADAY * AFTER			-0.044	-0.044
			(0.174)	(0.174)
Day-of-the-Week Indicators	Y	Y	Y	Y
Month Indicators	Y	N	Y	N
Time Trend	N	Y	N	Y
Num. Obs.	220	220	220	220
Adj. R-squared	0.222	0.221	0.212	0.212

Standard errors in parentheses. ** Significant at the 5% level. * Significant at the 10% level. The omitted intraday period is the first intraday period, from 6 a.m. to 10 a.m. The second intraday period is 10 a.m. to 2 p.m., the third intraday period is 2 p.m. to 6 p.m., and the fourth intraday period is 6 p.m. to 12 a.m. midnight.

TABLE VII

EFFECTS OF PRICE DISPERSION AND PRICE CHANGES ON SEARCH

Dep. Var.: ln(ReportCount)	(1)	(2)	(3)
SECOND_INTRADAY_PERIOD		-0.479**	-0.429*
		(0.064)	(0.101)
THIRD_INTRADAY_PERIOD		-0.325	-0.280*
		(0.147)	(0.093)
FOURTH_INTRADAY_PERIOD		-0.112	-0.020
		(0.046)	(0.010)
Δp_{t+1}^{+}	0.030		-0.021
	(0.020)		(0.013)
Δp_{t}^{+}	0.078**		0.020
	(0.006)		(0.008)
$\Delta p_{t-1:t-3}^{+}$	-0.006		0.002
	(0.009)		(0.006)
$\Delta p_{t-4:t-11}^{+}$	-0.007		-0.008
	(0.007)		(0.006)
Δp_{t+1}^{-}	0.015		0.007
	(0.034)		(0.022)
Δp_{t}^{-}	-0.014		-0.012
	(0.011)		(0.016)
$\Delta p_{t-1:p-3}^{-}$	0.027		0.005
	(0.014)		(0.006)
$\Delta p_{t-4:t-11}^{-}$	-0.013		-0.014
	(0.006)		(0.006)
City Indicators	Y	Y	Y
Day-of-the-Week Indicators	Y	Y	Y
Month Indicators	Y	Y	Y
Num. Obs.	650	660	650
Adj. R-squared	0.848	0.863	0.863

Standard errors in parentheses. ** Significant at the 5% level. * Significant at the 10% level. The omitted intraday period is the first intraday period, from 6 a.m. to 10 a.m. The second intraday period is 10 a.m. to 2 p.m., the third intraday period is 2 p.m. to 6 p.m., and the fourth intraday period is 6 p.m. to 12 a.m. midnight.