

Startup Search Costs*

David P. Byrne[†]

Nicolas de Roos[‡]

March 11, 2019

Abstract

Workhorse economic models used for studying the market impacts of search frictions assume constant search costs: individuals pay the same cost to obtain price information each time they search. This paper provides evidence on a new form of search costs: startup costs. Exploiting a natural experiment in retail gasoline, we document how a temporary, large exogenous shock to consumers' search incentives leads to a substantial, permanent increase in price search. A standard search model fails to explain such history-dependence in search, while it follows directly from a model with a one-time up-front cost to start searching.

JEL Classification: D83, L81

Keywords: Search, Price Dispersion, Price War, Retail Gasoline

*We recommend viewing this article in color. We thank Simon Loertscher, Eeva Mäuring, Wanda Mimra, Jose Moraga-Gonzalez, Philipp Schmidt-Dengler, Paul Scott, Matthew Shum, Michelle Sovinsky, Tom Wilkening, Nan Yang, and seminar participants from the University of Melbourne, University of Groningen, University of Vienna, ESMT Berlin, Mines ParisTech, University of Adelaide, University of Technology Sydney, NYU Stern School of Business, 2016 Asia-Pacific IO Conference, 2017 Monash IO Workshop, and 2018 Tinbergen Institute Empirical IO Workshop for helpful comments. The views expressed herein are solely those of the authors. Any errors or omissions are our own.

[†]Department of Economics, The University of Melbourne, Level 4 111 Barry Street, VIC 3010, Australia, e-mail: byrned@unimelb.edu.au.

[‡]School of Economics, The University of Sydney, Merewether Building H04, NSW 2006, Australia, e-mail: nico-las.deroos@sydney.edu.au.

1 Introduction

To purchase a flight, find a nice restaurant, or fill up one’s gas tank, consumers must set aside time and effort to search for the most suitable or affordable option. According to the theory of search, the trade-off between these search costs and the expected benefits from searching determines the set of prices and products considered by a consumer, and firms’ market power.¹ In applications of search models that study this trade-off, parsimony is important for tractability, and it is commonly assumed that the cost of search is independent of a consumer’s past experience searching for prices and products.

In this paper, we provide evidence on a new form of search costs that we call startup search costs. These costs represent a one-time, up-front cost and individual incurs the first time they engage in search behavior. After a startup search cost is sunk, their cost of search for future purchases is reduced. To establish the existence of such startup search costs, we leverage a unique dataset and natural experiment in retail gasoline.

Section 2 describes the dataset. The data come from an urban gasoline market that has an online price clearinghouse which consumers can use to become informed about the retail price distribution across stations day-to-day. We have access to the universe of station-level gasoline prices from 2014-2016, which we match to daily website usage from the price clearinghouse. Our ability to perfectly measure price levels and dispersion, and directly measure market-level search behavior from a price clearinghouse in a homogeneous product market, makes this context particularly well-suited for studying the dynamics of search and price dispersion.²

In Section 3, we describe the natural experiment which occurs mid-sample. We then exploit the experiment to identify the causal impact of a temporary, large exogenous shock to consumers’ search incentives on search intensity in the short and long-run.

The natural experiment stems from a price war that disrupts a stable coordinated pricing equilibrium. In a separate study, [Byrne and de Roos \(2019\)](#), we document *five years* of stable and highly coordinated pricing in the market from March 2010 to April 2015, which is consistent with tacit collusion. However, in May 2015, a price war occurs and price coordination breaks down. During a three-week conflict period, retail price dispersion spikes and the daily gains from search grow by 100% relative to their pre-war levels. Firms eventually resolve the war and return to stable coordinated pricing, and price dispersion and search incentives return to their baseline (pre-war) levels. Based on pre-war retail price and search dynamics in our dataset,

¹[Baye, Morgan, and Scholten \(2006\)](#) overview an extensive and influential literature on search models, which date back to [Stigler \(1961\)](#).

²See, for example, [de Los Santos, Hortascu, and Wildenbeest \(2012\)](#) and [Koulayev \(2014\)](#) for analyses of search models using web search data. See [Eckert \(2013\)](#) for an overview of an extensive literature on search and price dispersion in retail gasoline.

and anecdotal evidence of a retail ownership change that precipitated the war, we argue that the price war is exogenous to search behavior on the price clearinghouse.

We study search intensity on the price clearinghouse one year before and after this temporary shock to price dispersion and search incentives. Our main empirical result is that daily search intensity permanently rises by 70% following the shock. We show that this rise in search intensity comes from an increase in the number of unique searchers using the clearinghouse, and not an increase in average search intensity among previous searchers. This finding of history-dependence in search behavior points to new consumers experimenting with the clearinghouse for a first time in response to the substantial price uncertainty created by the price war. Having tried it once, consumers continue to use the clearinghouse thereafter, even after search incentives return to pre-war levels.

Section 4 moves beyond these reduced-form empirics and formalizes our notion of startup search costs, estimates their magnitude, and shows how ignoring them leads to model misspecification. Specifically, we develop and estimate a non-sequential search model that decomposes search costs into a *startup* and a *recurrent* component.³ Exploiting the natural experiment, we estimate the relative magnitudes of startup search costs and recurrent search costs. Our simple search model rationalizes the permanent rise in search intensity from a temporary shock to search incentives. By contrast, a standard non-sequential search model is unable to rationalize the persistent post-war increase in search that we find.

We conclude in Section 5. Our findings suggest that the omission of startup search costs leads to qualitatively different inferences regarding search behavior. Moreover, our study illustrates the challenges faced by policymakers when designing tools to enhance price transparency when consumers face startup search costs: the price clearinghouse we study had been in existence for 15 years by May 2015, more than 90% of the market were aware of its existence, yet the price war shock still led to a 70% rise in the volume of search. Finally, we highlight avenues for future research. While we provide evidence on the existence of startup search costs and measure their magnitude, further work is required to elicit their microeconomic foundations.⁴ We close by discussing deeper possible mechanisms for the novel history-dependent search behavior that we uncover.

³Under non-sequential search, the entire vector of prices is revealed to a consumer that incurs a search cost. This matches the market environment in which a comprehensive price clearinghouse operates; see Ellison (2016).

⁴In this way, our startup search cost estimates are reduced forms like previous search cost estimates based on sequential and non-sequential search models (e.g., Hong and Shum, 2006; de Los Santos et al., 2012; Koulayev, 2014). Only recently have studies emerged that identify deeper microfoundations of search costs including spatial frictions (Pennerstorfer et al., 2018; Startz, 2018; Buchholz, 2018), salience and price obfuscation (Blake et al., 2018), and the value of time while searching (Seiler and Pinna, 2017).

2 Context and data

Our research context is Perth, Australia, a city with approximately 2 million people. Perth, like many urban gasoline markets worldwide, has a concentrated retail gasoline market. Four major firms dominate the market: BP, Caltex, Coles and Woolworths. The former two firms are vertically integrated oil majors, while the latter two firms are major supermarket chains that also sell gasoline. All four firms either directly or indirectly control retail prices day-to-day at their large station networks. Combined, their station networks account for approximately 75% of stations in the market. All other stations are operated by independent retailers.

A key aspect of the market is a price transparency program called Fuelwatch. The program was introduced on January 3, 2001 by Western Australia's state government. By law, before 2pm each day firms must submit CSV files to the state government that contain tomorrow's station-level retail prices. The next day at 6am when stations open, they are required to post the prices that were submitted at 2pm the previous day. Prices are then fixed for 24 hours. From our conversations with the government, compliance with the program is near perfect.

With the Fuelwatch price data, before 2pm each day the government posts online today's prices for all stations in the market. This helps customers engage in cross-sectional price search. After a data verification check, at 2:30pm the government further posts tomorrow's price for every station in the market. This helps customers engage in inter-temporal price search. Figure 1 depicts the Fuelwatch price clearinghouse at www.fuelwatch.gov.au. Survey evidence from the Western Australian Government indicates that more than 90% of Perth households are aware of the Fuelwatch price clearinghouse.⁵

2.1 Data

The price clearinghouse generates a uniquely rich dataset for studying retail price search and price dispersion. From the clearinghouse, we have access to the universe of retail prices from 2001 to present day. This allows us to perfectly measure daily price levels and dispersion. We match these price data to the daily terminal gate price (TGP) for gasoline, which is the local spot price for wholesale gasoline. These spot prices include a margin for upstream suppliers of gasoline. The difference between the retail price and TGP provides an estimate of the retail margin. It is an appropriate estimate for studying the evolution of margins over time because

⁵For example, in 2008 the Deputy Prices Commissioner of the Western Australian Department of Consumer and Employment Protection stated that "The last survey we did showed that 94 per cent of people had heard of FuelWatch." See Commonwealth, *Parliamentary Debates*, Senate Standing Committee on Economics, 16 July 2008, (Mr. Rayner), p. E 5.

Figure 1: Fuelwatch Price Clearinghouse

FuelWatch
Consumer Protection
Subscribe Forgot Password

My FuelWatch Login
Email Address
Password
Go

Home About FuelWatch My FuelWatch Price Search FuelWatch News Fuel Information For Industry

FuelWatch Quick Search - Results
Product: ULP
Metro Region: Any Metro
Suburbs: PERTH (including surrounding suburbs)
Brands: Any Brand
Country Region: None
Date: Today and tomorrow
Refine Search New Search

Best prices available from 6:00am for today and tomorrow

Today	Tomorrow	Product	Brand	Name <small>Mouse over Name for details</small>	Address	Suburb/Town	Map
92.1	92.1	ULP	Coles Express	Coles Express Mount Lawley	Cnr Walcott St & Fitzgerald St	NORTH PERTH	
93.1	93.1	ULP	Caltex	Caltex StarMart Highgate	Cnr Beaufort St & Bulwer St	HIGHGATE	
93.1	93.1	ULP	Caltex	Caltex StarMart East Perth	157 Lord St	PERTH	
93.1	93.1	ULP	Caltex	Caltex Wellington Street	141 Wellington St	PERTH	
99.9	99.9	ULP	United	United Mt Lawley	791 Beaufort St	MT LAWLEY	
99.9	99.9	ULP	Caltex	Caltex StarShop Mt Lawley	810 Beaufort St	MT LAWLEY	
104.9	92.1	ULP	Puma	Puma First Avenue Mt Lawley	81 Guildford Rd	MT LAWLEY	
106.7	103.7	ULP	United	United Northbridge	31 Fitzgerald St	NORTHBRIDGE	
107.5	107.5	ULP	Vibe	Vibe Charles St	427 Charles St	NORTH PERTH	
107.9	107.9	ULP	BP	BP Connect East Perth	Cnr East Pde & Brown St	EAST PERTH	
107.9	107.9	ULP	BP	BP Connect North Perth	Cnr Scarborough Beach Rd & Charles St	NORTH PERTH	
109.4	109.4	ULP	Gull	Gull East Perth	Cnr Pier St & Brisbane St	PERTH	
109.9	99.9	ULP	Coles Express	Coles Express Highgate	1-5 Guildford Rd	MT LAWLEY	
109.9	99.9	ULP	Coles Express	Coles Express Perth	480 William St	PERTH	
109.9	99.9	ULP	Coles Express	Coles Express West Perth	Cnr Wellington St & Thomas St	WEST PERTH	
125.9	125.9	ULP	Wesco	Wesco North Perth	41 Angove St	NORTH PERTH	

16 Prices found

the TGP is the main time-varying component of stations' wholesale fuel costs.⁶

Moreover, the state government provided us access to daily website usage data from the Fuelwatch website. More specifically, we were provided daily data on the total number of website hits on the Fuelwatch website, and the total number of unique visitors to the Fuelwatch website month-to-month. These search data, combined with the universe of station-level price data, permit a direct examination of how retail search behavior responds to changes in price levels and dispersion at the market level over time.⁷

Our sample period spans two-years, from July 1, 2014 to June 30, 2016. Over this period, average station-level retail prices and margins in terms of cents per liter (cpl) are 128.6 cpl and 10.3 cpl, respectively. The average number of Fuelwatch website visits each day is 19699, and there are 208184 unique visitors to the website each month on average. Table 1 presents summary statistics from our dataset.

⁶Other time-invariant parts of marginal costs include quantity discounts, shipping costs, wharfage fees, and insurance costs.

⁷Daily data on the total number of unique visitors is unavailable due to data privacy concerns from the state government. Similarly, daily data on website hits at the individual level or by disaggregated census blocks are not available because of privacy concerns.

Table 1: Summary Statistics

	Obs.	Mean	Std. Dev	Min	Max
<i>Price data</i>					
Station Retail Price	146931	128.64	13.67	92.9	169
Market Terminal Gate Price	731	118.3	12.6	95.8	147.0
Station Retail Margin	146931	10.3	6.8	-17.0	35.7
<i>Search data</i>					
Daily website visits	731	19699	6691	8243	53517
Monthly unique website visitors	24	208184	47767	129630	276737

Notes: Sample period is July 1, 2014 to June 30, 2016.

2.2 Price cycles and search cycles

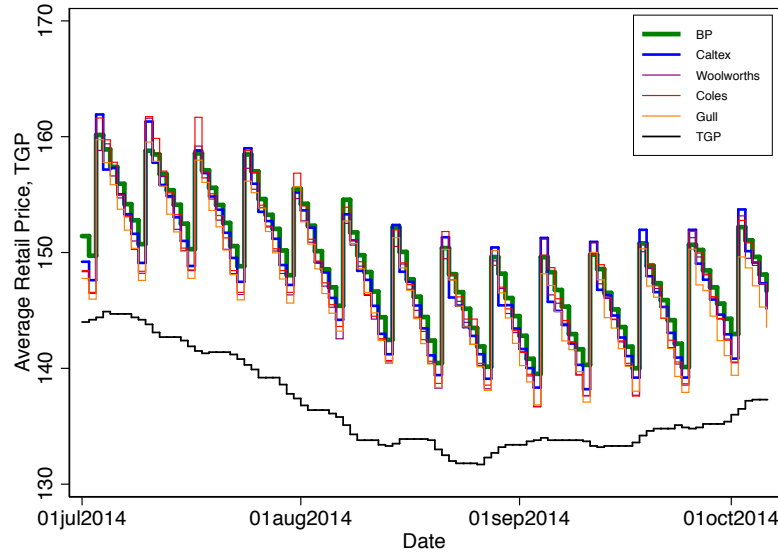
Figure 2 depicts cyclical patterns in retail price levels, price dispersion, and search that persist over the entire sample period. Panel A, which plots average daily retail prices by retailer and the daily TGP, highlights cycles in price levels. Over the time period shown in the figure, every Thursday, retail prices jump by approximately 10% of the average market price. Between price jumps, most retailers cut prices by 2 cpl each day until the next price jump day occurs. Intertemporal search incentives are therefore heightened on Wednesdays in anticipation of Thursday price jumps.

Panel B plots corresponding daily market price dispersion, as measured by the daily standard deviation of retail prices across stations. For reference, we overlay daily price dispersion with the price cycles from panel A in greyscale. The figure highlights spikes in price dispersion on Thursdays, as the average of the standard deviation of retail prices rises from 4.97 at the bottom of the cycle, to 6.77 cpl at the top of the cycle. This large rise in price dispersion reflects a noisy coordination process as market prices transition from the bottom to the top of the cycle. Cross-sectional search incentives thus tend to rise substantially on Thursdays.

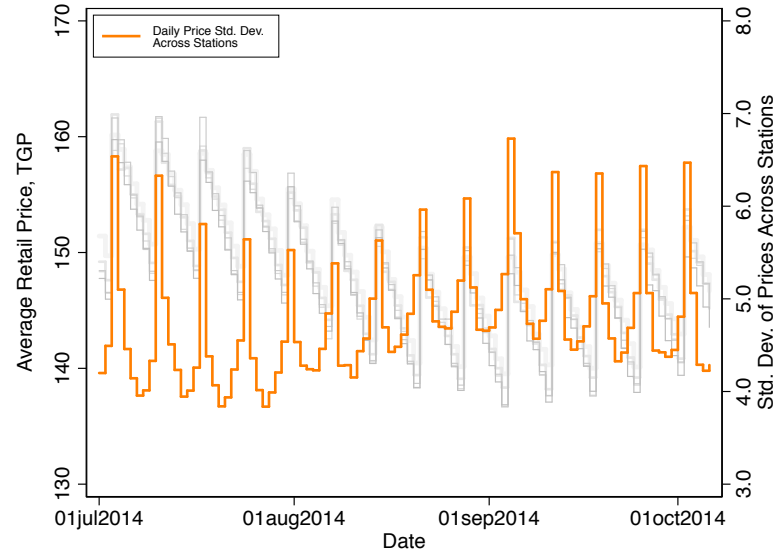
Panel C of Figure 2 plots daily search intensity on the Fuelwatch website, as measured by the number of website hits from computers and mobile devices. Like prices, search exhibits cycles, whereby search intensity jumps on Wednesdays and Thursdays. On the day before and day of price jumps, on average the website receives 27049 and 19531 visits, respectively. On all other days, the website on average receives 18263 visits. In sum, search intensity rises with cross-sectional and inter-temporal search incentives just before and during price jumps.

Figure 2: Price and Search Cycles Cycles

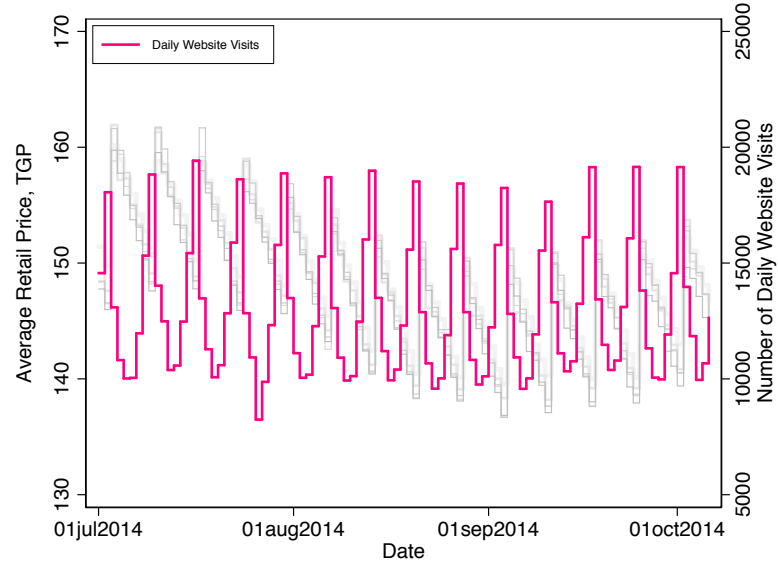
Panel A: Daily Retail Price Levels



Panel B: Daily Retail Price Dispersion



Panel C: Daily Number of Website Visits



3 Dynamics of price dispersion and search

In this section, we examine the coevolution of retail prices and search over the entire two-year sample period. We document the break out of a price war among the retailers who were previously coordinating on the timing and magnitude of price jumps and cuts. We describe the impact of the price war on price levels, dispersion, and margins. We then show that daily search intensity permanently rises by 70% following the price war, and argue that this response of search to a temporary, war-induced shock to search incentives is causal.

3.1 Price War

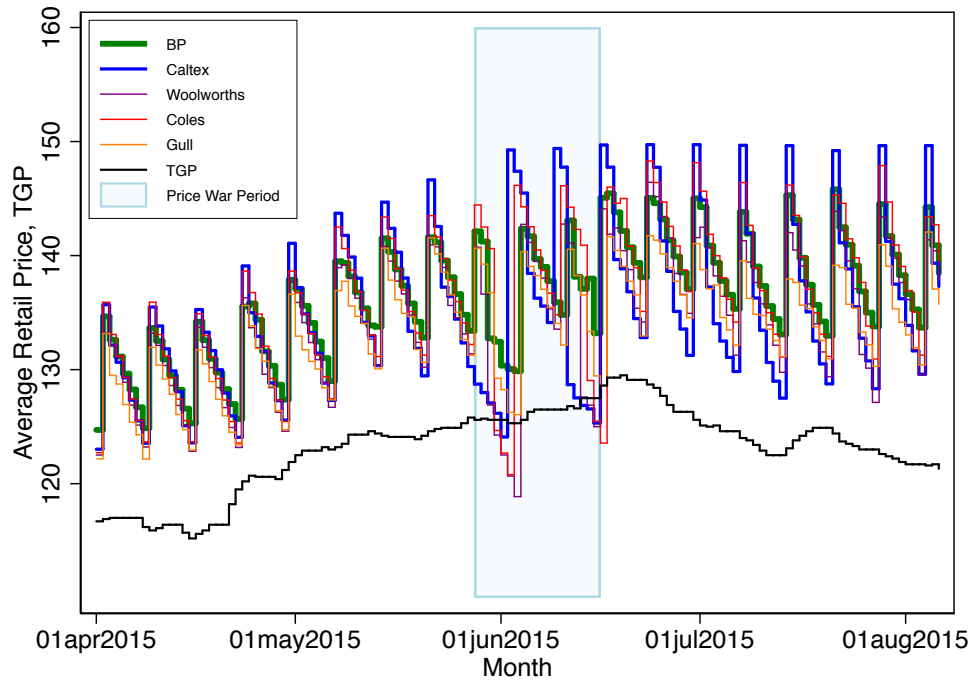
As alluded to above, the price cycles in Figure 2 are stable and regular. As we document in Byrne and de Roos (2019), between March 2010 and April 2015, firms coordinate on price cycles using two simple focal points: Thursday price jumps and 2 cpl price cuts on days of the week between jumps. Over this five year period *every* market price jump occurs on Thursday. The result is a tightly coordinated price cycle that is consistent with tacit collusion.

However, after five years of stability, Caltex breaks from this pricing strategy in May 2015. This is depicted in Figure 3, which plots average daily retail prices for the four major retailers and an independent retailer, Gull. The figure highlights the disruption Caltex's defection creates, particularly with the timing of price jumps. In particular, in the last week of May, Caltex defects from Thursday price jumps and instead engages in Tuesday jumps the following week in the first week of June. After three weeks of turmoil and an uncoordinated cycle, the other major retailers along with Gull transition to Tuesday price jumps, re-establishing a coordinated price cycle. Tuesday price jumps are subsequently stable from June 15 to present day (Australian Competition and Consumer Commission, 2017).

Exogeneity of the Price War to Search. We argue that the price war and its corresponding effects on price levels and dispersion is exogenous to search behavior with the price clearing-house. We believe this for three reasons. First, as mentioned, the coordinated pricing structure is stable for five years prior to the May 2015 price war, and the price war occurred without warning. Indeed, media releases from the Western Australian government and local major news outlets highlight a high degree of unpredictability of prices during the price war. For example, a manager at the Western Australian Royal Automobile Club describes the public's surprise in June 2017 over the price war and transition to Tuesday price jumps as follows:⁸

⁸The quote appeared on July 6, 2015 in an Australian Broadcasting Corporation article entitled "Monday cheapest day to buy fuel in Perth in change to long-running petrol cycle", accessed at <http://www.abc.net.au/news/2015-07-06/monday-now-cheapest-for-fuel-in-perth/6598290>.

Figure 3: 2015 Price War



“We don’t have a full understanding of why the cycle has changed ... and we want to understand why this is happening. It has changed after a period of certainty and we don’t know what the future looks like”

Second, based on our conversations with the Western Australian Government, the only shock in the market that can potentially be linked to the price war is a major change in ownership at Caltex. In March 2015, the Chevron Corporation sold off its 50% share in Caltex Petroleum Australia Pty. Ltd. to Australian shareholders; after this sale, Caltex becomes 100% owned by shareholders.⁹ This supply-side shock in ownership may have led to a change in management and pricing tactics, and hence the price war. However, the ownership change is unrelated to local demand or online price search in Perth.

Finally, as we show momentarily in Figure 4 below, search behavior on the Fuelwatch price clearinghouse is stable for the entire year prior to the price war. There is no evidence to suggest that changes in online search precipitates the price war.

Given these empirics and anecdotal evidence, we assume that the price war and related changes in price dispersion are exogenous to search behavior on the price clearinghouse. We therefore interpret the empirical relationship between search and price dispersion as causal

⁹See, for example, the Australian Competition and Consumer Commission “Report on the Australian petroleum market”, March Quarter 2015.

and reflecting demand-side search behavior. We further pursue this interpretation in estimating a search model in Section 4 below.

3.2 Evolution of search and price dispersion

Figure 4 presents time series for daily average retail price levels (panel A), margins (panel B), price dispersion (panel C), and search (panel D) from July 1, 2014 to July 1, 2016. In each panel, we plot the raw daily time series in greyscale, and the weekly average of the daily series in color, which more clearly depicts trends. The shaded vertical area in the middle of each figure highlights the price war period.

In panel A we see that price levels primarily trend with wholesale costs over time. Panel B highlights cyclical daily margins that arise because of the price cycle. At weekly frequencies, margins hover at around 10 cpl, and average 7 cpl during the price war. There are no other discernable trends in margins before or after the price war.

Panel C illustrates how price dispersion, as measured by the daily standard deviation of prices, rises by more than 100% around the price war. During this period, search incentives are higher than they have been in the past five years. Dispersion drops from its peak immediately after the price war is resolved, and gradually returns to pre-war levels six months after the price war in January 2016.¹⁰ That is, the price war creates a large, temporary exogenous shock to search incentives in the market.

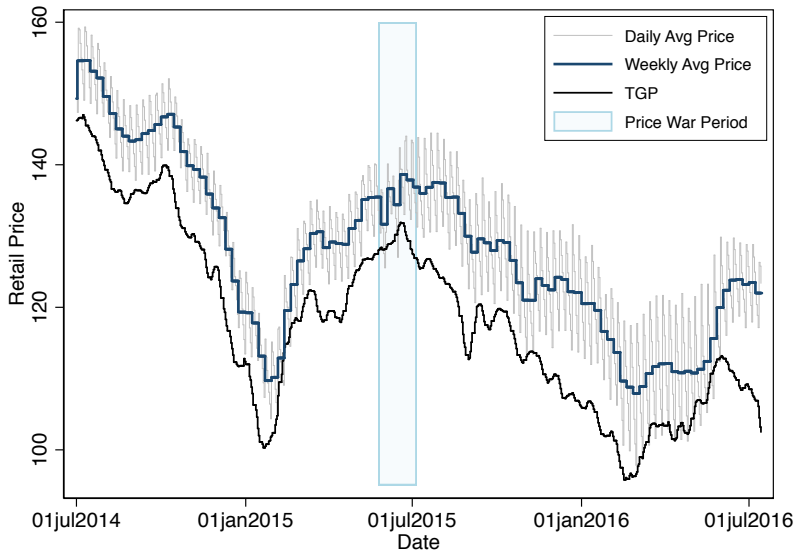
How does search respond to these price fluctuations? Panel D of Figure 4, which we view as the paper’s central result, provides the answer: despite the temporary shock to search incentives, we find an immediate and permanent increase in search intensity with the price clearinghouse. This increase in daily search is also large: it rises from an average of 14097 visits before the shock to 24461 visits after the shock, a 70% increase. Importantly, this jump is *not* driven by automated emails or text messages from Fuelwatch that might cue search behavior; we have confirmed with Fuelwatch that such messages are not sent to consumers. The shift in panel D reflects new and permanent active effort in using the price clearinghouse following the shock.

This substantial and permanent rise in search intensity could reflect a rise in search intensity among existing searchers, or the emergence of new searchers, or both. Figure 5 provides evidence that strongly suggests it is driven by new searchers. With the left axis we plot the number of unique visitors to the Fuelwatch price clearinghouse month-to-month. We find the average number of unique visitors permanently rises from 151,677 visitors pre-war to 239,959

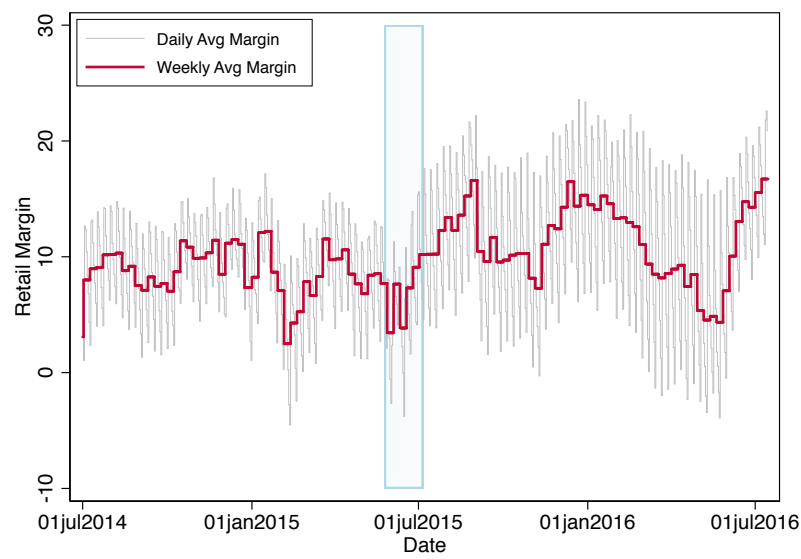
¹⁰While the transition to Tuesday price jumps is completed after the three week price war, cross-sectional price dispersion remains elevated on price jump days for several months after the price war. This dispersion reflects larger price jumps by Caltex relative to its rivals. By November 2015, firms are again able to coordinate on the size of price jumps, and price dispersion returns to baseline.

Figure 4: Prices, Margins and Search Before and After the Price War

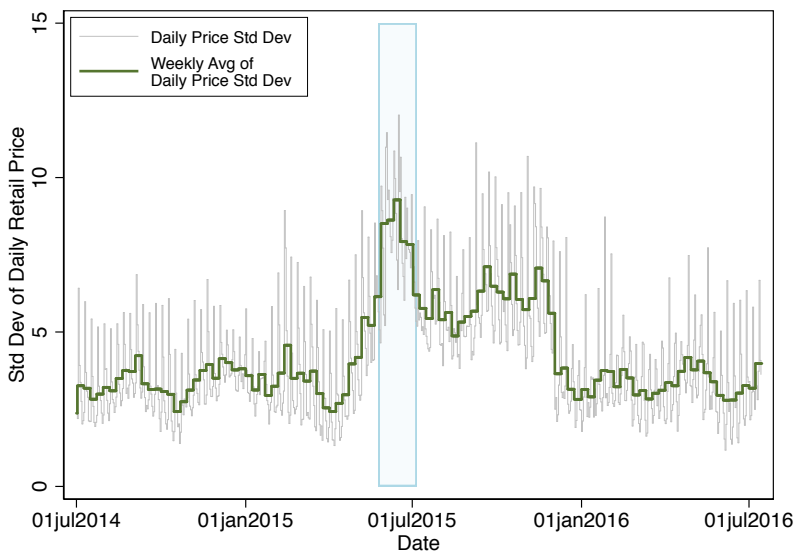
Panel A: Price Levels



Panel B: Margins



Panel C: Price Dispersion



Panel D: Search

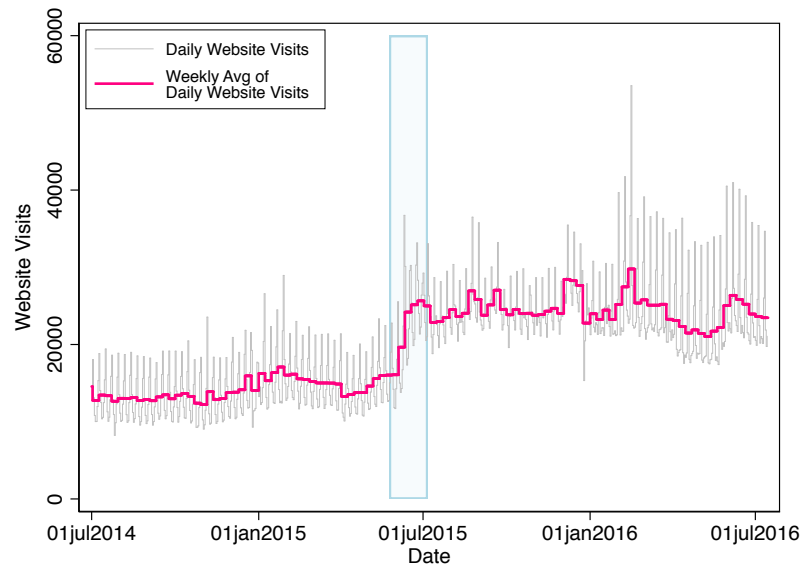
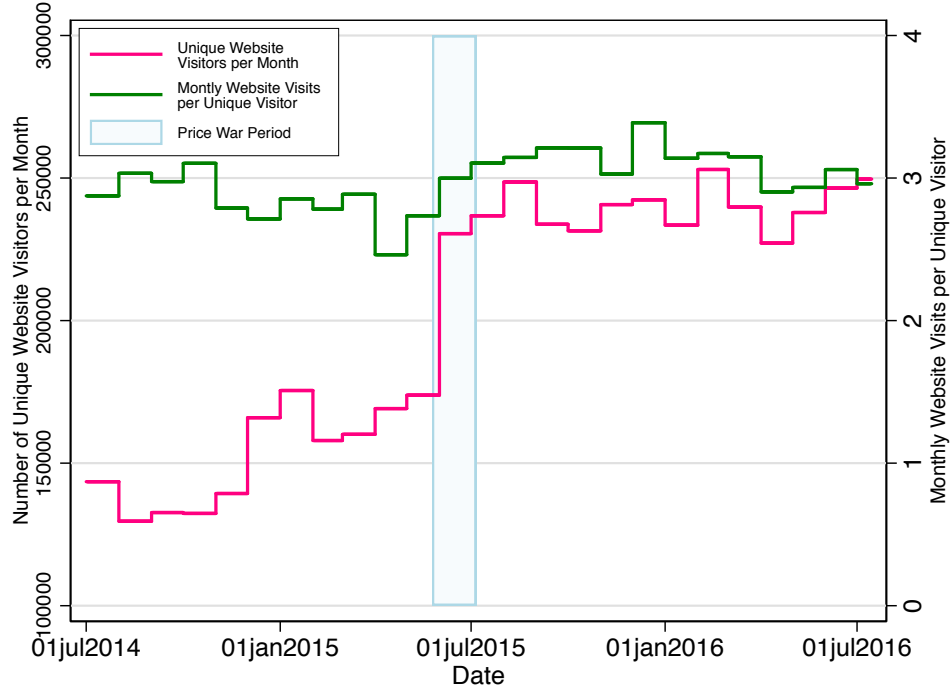


Figure 5: Unique Searchers and Search Intensity Before and After the Price War



visitors post-war, a 60% increase in the number of unique searchers. With the figure’s right axis, we plot the number of visits to the Fuelwatch price clearinghouse per unique visitor month-to-month. This remains stable around three visits per month before and after the price war. That is, search intensity per searcher does not appear to dramatically rise after the price war.

These empirics collectively present a challenge for conventional search models. If searchers are myopic and face search costs that are unchanged irrespective of past search behavior, then search levels should return to their baseline levels as search incentives return to their baseline levels over time. This is clearly not the case in Figures 4 and 5. Past search experience appears to be important for future search behavior. That is, we find history dependence in search. These results suggest that the first time a consumer engages in search, “startup” search costs are potentially high. Conditional on sinking these costs, the persistence in search levels among new searchers following the shock indicates that recurrent search costs from using the price clearinghouse are small.

4 Estimating startup search costs

In this section, we formalize these ideas by estimating a search model that incorporates startup search costs in an otherwise standard model of non-sequential price search. We introduce the model in Section 4.1, discuss estimation and identification in Section 4.2, and present our results in Section 4.3.

Our goal is to provide indicative estimates of the relative magnitudes of startup and recurrent search costs in a parsimonious model. We specify non-sequential search on the price clearinghouse because households obtain information on the entire distribution of prices when they search on it. We also presume that consumers are unsophisticated in the sense that they consider only the current-period benefits of a decision to use the search platform. A sophisticated consumer, aware of her startup search costs, would weigh the expected net present value of future benefits of learning to use the search technology, leading to higher estimates of startup search costs.¹¹ Our specification also abstracts from the incentive for intertemporal search. In [Byrne and de Roos \(2017\)](#), we find evidence for intertemporal search in this market, which contributes to within-week variation in search. By focusing on cross-sectional search, our model is better suited to identifying longer term trends in search rather than within-week fluctuations.

Our model setup reflects the data we work with. We have market-level search data, and therefore we are unable to identify individual-level heterogeneity in the model’s parameters.¹² Moreover, as with virtually all research on gasoline demand, we do not have access to high-frequency data on quantities of gasoline purchased and consumers’ fuel tank inventories.¹³

We focus strictly on the demand-side of the market for two reasons. First, as argued in Section 3, daily price changes are plausibly exogenous to online search. Given that we have a direct measure of search intensity, we can use the demand side of a search model alone to identify the search costs from using the price clearinghouse. Second, in [Byrne and de Roos \(2019\)](#) we argue that firm behavior is consistent with tacit collusion over our sample period. Therefore, we cannot use standard static first-order conditions to model pricing and build supply-side moments, as in [Hong and Shum \(2006\)](#) or [Wildenbeest \(2011\)](#), to help identify search costs. A dynamic model of the supply side of the market is beyond the scope of the current paper.

¹¹In Appendix A we show that when consumers are sophisticated, inferred startup costs are higher for more patient consumers, while inferred recurrent search costs are unaffected.

¹²See, for example, [Koulayev \(2014\)](#) for an application that studies online price search using individual-level search behavior from an online hotel price search website. As mentioned in Section 2, we requested such data for Fuelwatch from the Western Australian Government. Unfortunately, they are prohibited from providing such search data at the individual level or disaggregated into local areas.

¹³[Levin, Lewis, and Wolak \(2017\)](#) provide the first ever published research on daily market-level gasoline demand behavior. With their unique data, they are able to distinguish between the binary decision to purchase gasoline and how much gasoline to purchase conditional on deciding to purchase. Their reduced-form study abstracts from search frictions and inventories.

4.1 Model

Consumer i 's indirect utility on date t if she purchases κ_i liters of petrol is given by

$$u(s_{it}) = \begin{cases} \bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - s_{it} & \text{if she searches} \\ \bar{u} - \kappa_i \bar{p}_{it} & \text{otherwise,} \end{cases} \quad (1)$$

where \mathbf{p}_{it} is the vector of prices available to consumer i at date t , and s_{it} is the current cost of search for consumer i . As we discuss below, the set of stations available to consumer i is based on geographic proximity to her home address, L_i . This formulation assumes that if consumer i on date t engages in price search on the clearinghouse, she becomes fully informed about the price distribution and pays the minimum price in her local choice set, $\min\{\mathbf{p}_{it}\}$. If she does not search, she pays the average price in her local choice set, \bar{p}_{it} . Her gains from searching in period t are

$$g_{it} = (\bar{p}_{it} - \min\{\mathbf{p}_{it}\}) \times \kappa_i.$$

Consumer i 's search costs in period t are given by

$$\begin{aligned} s_{it} &= f_i \times (1 - w_{it}) + c_i, \\ w_{it} &= 1\{\text{consumer } i \text{ has searched before date } t\}, \end{aligned}$$

where $1\{\cdot\}$ is an indicator function, and f_i and c_i are consumer i 's startup and recurrent search costs, respectively. Recall that the Fuelwatch price clearinghouse allows consumers searching after 2:30pm to discover prices for today and tomorrow. To account for this, consumer i considers the gains from search in periods t and $t + 1$ when making her search decision on date t :

$$y_{it} = 1\{\max\{g_{it}, g_{it+1}\} > s_{it}\}.$$

Define the type of consumer i as $\tau_i = (f_i, c_i, L_i)$. We assume that f_i , c_i , and L_i are independent, and that the search costs f_i and c_i are drawn from gamma distributions (Hong and Shum, 2006) with shape parameters μ_f and μ_c and scale parameters σ_f and σ_c , respectively.¹⁴ Consumer locations L_i are drawn from the empirical distribution of population locations in the market, as described below. We collect the search cost parameters into $\theta = [\mu_f, \sigma_f, \mu_c, \sigma_c]'$. Let P be the distribution of consumer types. The predicted share of consumers engaging in online

¹⁴We have estimated the model under alternative functional form assumptions, with no qualitative differences in reported results. In particular, we have estimated the model under the assumption that f_i and c_i are drawn from the Log Normal distribution, both under the assumption that startup and recurrent search costs are independent, and allowing for correlation between them.

price search on date t , $q_t(\theta)$, is obtained by integrating over P :

$$q_t(\theta) = \int y_{it}(\tau_i) dP(\tau_i; \theta) = \int y_{it}(\tau_i) dP_f(f_i; \mu_f, \sigma_f) dP_c(c_i; \mu_c, \sigma_c) dP_L(L_i). \quad (2)$$

4.2 Estimation and identification

We estimate the model using a Simulated Minimum Distance estimator that compares the share of searchers predicted by the model, $q_t(\theta)$, to its empirical analogue, \hat{q}_t , computed as

$$\hat{q}_t = \frac{n_t}{Q},$$

where n_t is the number of Fuelwatch website hits on date t , and the market size Q represents the number of consumers considering a gasoline purchase each day.¹⁵

Computing \hat{q}_t requires us to make an assumption regarding the market size, Q . We compute this as $Q = (0.80 \times 1,576,000)/7$, which assumes that 80% of the population in Perth aged between 15 and 79 years plans to fill up their car once every week, and does so uniformly by day of the week. We calibrate the size of a gasoline purchase to $\kappa_i = 50$ liters for all consumers. The most popular car in Australia is the Toyota Corolla, which has a 55 liter tank. According to this calibration, each representative consumer fills up their Toyota Corolla when it is almost empty.¹⁶

To construct the distribution of consumer locations P_L , which determines the distribution of gains from search, we partition the region of Greater Perth into local districts as classified by the Australian Bureau of Statistics.¹⁷ For each district, we obtain driving age (18-79 years) population and the location of the centroid of the district from the 2011 Census. By matching this to the location of each station, we obtain the set of stations within a 5km radius.¹⁸ Average

¹⁵By fixing the market size over time, we implicitly introduce an outside good into the model. Because some consumers are aware of the cyclical nature of prices in this market, there is a cycle in sales volumes over the week. See, for example, [Australian Competition and Consumer Commission \(2014\)](#) on the existence of a demand cycle in the Perth market. Fixing the market size amounts to assuming that, each day, the same volume of consumers considers both whether to purchase gasoline and whether to use the search platform.

¹⁶Data sources for our calculations are as follows. From Australian Bureau of Statistics Table 3235.0, "Population by Age and Sex, Regions of Australia", there were 1,576,479 people aged between 15 and 79 years in the greater Perth area on June 30, 2014. The [Australian Competition and Consumer Commission \(2007\)](#) report into the Australian petrol market commissioned a survey of 775 motorists in Australia, finding that 26% purchase more than once per week, 50% purchase once per week, 20% purchase every 2 weeks, and 4% purchase less than every 2 weeks. According to the Federal Chamber of Automotive Industries (FCAI), the most popular car in Australia in 2014 was the Toyota Corolla (see: <https://www.drive.com.au/motor-news/the-10-most-popular-cars-of-2014-20150105-12ihkp>). The fuel tank capacity for a Toyota Corolla is 55 liters (see <http://www.toyota.com.au/corolla/specifications/ascent-sedan-manual>).

¹⁷We use the finest classification available from the 2011 Census, known as Statistical Areas, Level 1. Of the 3789 Statistical Areas in the Greater Perth region, the average district has 318.7 people (s.d. 128.3).

¹⁸Our main qualitative findings are unaffected by the choice of search radius. However, as we might expect, in-

and minimum prices for each district are defined with reference to this local set of stations. To construct our simulated minimum distance estimator, we then assign each consumer randomly to a district according to weights based on driving age population. By calculating local search gains in this manner, we abstract from commuting patterns.¹⁹

Let $e_t(\theta) = q_t(\theta) - \hat{q}_t$ be the difference between the model's prediction and the fraction of searchers in the data at date t , and let $e(\theta)$ be the $T \times 1$ vector of prediction errors, where T is the number of dates in the sample. We estimate θ by minimizing the objective function²⁰

$$\hat{\theta} = \arg \min_{\theta} G(\theta) = e(\theta)'e(\theta).$$

For a given value of θ , we compute $G(\theta)$ by forward simulating the search shares, $q_t(\theta)$, for each sample date $t = 1, \dots, T$. To simulate search shares, we must keep track of each simulated consumer's history of search. Define the set of "active" consumers at the start of period t as

$$A_t = \{i : w_{i,t-1} = 1\}.$$

For the forward simulation, we must initialise the set of active consumers at the start of the sample period, A_1 . We use the maximum search intensity in the two weeks prior to the start of the estimation sample to define A_1 . In particular, taking T_0 to be the two week pre-sample, we set the size of A_1 equal to $\max_{t \in T_0} n_t$. We then assign active status randomly across consumers. With this method, we estimate that 10% of consumers in the market have already sunk their startup search costs at the start of the sample period.

Identification. The distribution of recurrent search costs is identified by periods in which the benefits of search are unremarkable. Thus, variation in aggregate search activity associated with variation in search gains, when such gains are moderate, identifies the parameters μ_c and σ_c . By contrast, the startup search cost parameters μ_f and σ_f are identified by the responsiveness of aggregate search to unprecedented gains from search arising from the price war. Finally, over the sample period, aggregate search varies between 5% and 30% of consumers, which we show in Figure 7 below. Therefore, estimation is well-suited for identifying the search cost distribution corresponding to this range of search intensities, but not necessarily the entire search cost distribution.

creasing the radius of search raises the value of information, and leads us to infer higher search costs. See Appendix B, where we also estimate the model using a 2km and a 10km search radius.

¹⁹See Pennerstorfer et al. (2018) for an analysis of commuting routes and search incentives.

²⁰We use the following optimization procedure. We first calculate the objective function $G(\theta)$ for a grid of parameters. We then calculate the minimum of $G(\theta)$ across the grid, and select all parameter vectors with a value of $G(\theta)$ that is within 10% of the minimum. We then perform a Nelder-Mead optimization for each of these parameter vectors. For our final estimates, we take 100,000 draws from P to evaluate the integral in equation (2).

Table 2: Search Model Estimation Results

	With startup Search Costs (1)	Without startup Search Costs (2)
Recurrent search cost distribution		
μ_c	0.202 (0.016)	0.611 (0.007)
σ_c	40.769 (3.039)	160.257 (5.475)
Startup search cost distribution		
μ_f	5.070 (0.086)	
σ_f	4.228 (0.203)	
Objective function, $G(\hat{\theta})$	0.356	1.061

Notes: Robust standard errors are in parentheses. The number of observations is $T = 731$ dates. All calculations assume consumers purchase 50 liters of gasoline.

4.3 Results

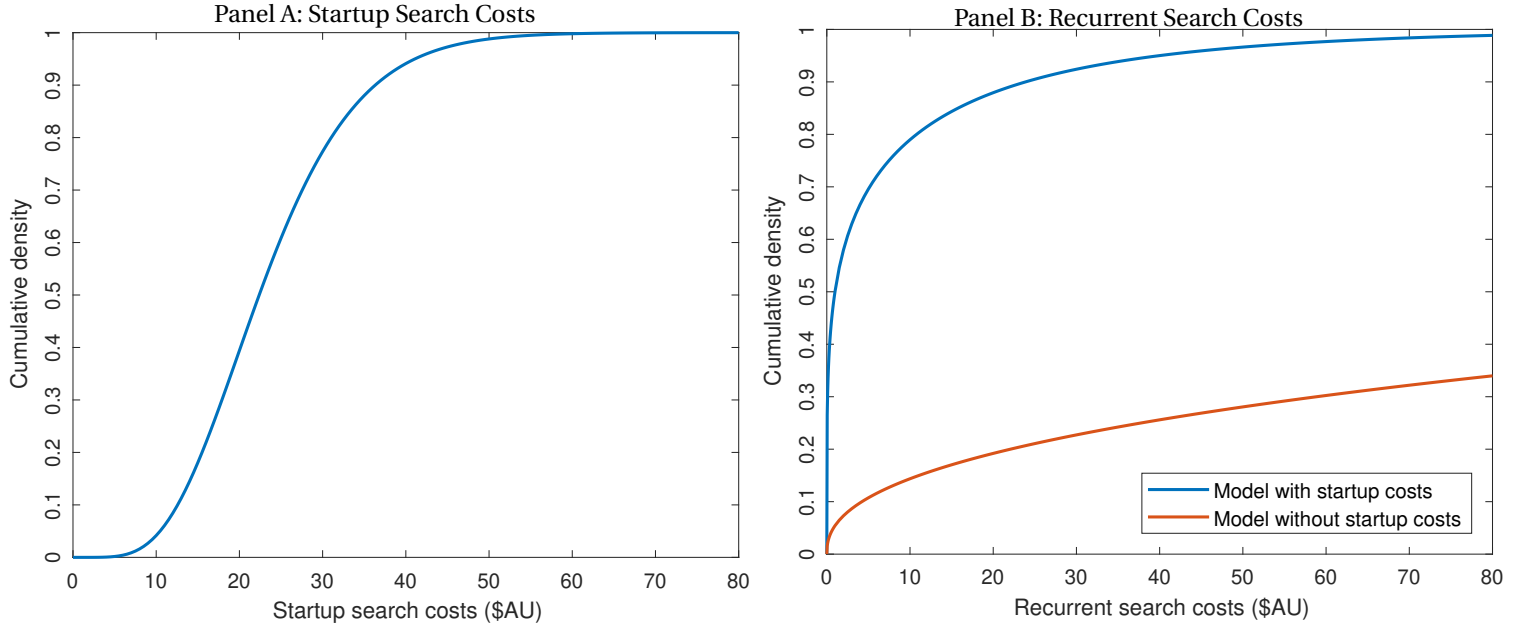
Estimation results are presented in Table 2. Column (1) contains estimates for the full model, and column (2) contains estimates for a constrained model in which there are no startup search costs. Panels (A) and (B) of Figure 6 present the cumulative density functions of the startup and recurrent search cost distributions based on the point estimates of the model.

To illustrate the qualitative difference between the models, consider the 20th percentile consumer in the search cost distributions. For the model with startup search costs, the 20th percentile startup costs are \$13.30/day. Conditional on having sunk this startup cost, the 20th percentile recurrent search cost is \$0.01/day.²¹ In the model without startup search costs, estimated recurrent search costs are an order of magnitude greater. In this model, the 20th percentile consumer has recurrent search cost of \$9.94. This conforms with intuition, as these recurrent search cost estimates are driven by both startup search costs and recurrent search costs.

Turning to model fit, from Table 2 we find the econometric objective function for the full model of $G(\hat{\theta}) = 0.356$ is substantially lower than the model without startup search costs, where $G(\hat{\theta}) = 1.061$. Figure 7 further describes the implications for model fit from accounting for

²¹Note that, because startup and nonsequential search costs are independently drawn, one cannot simply add the 20th percentile of each distribution to obtain the 20th percentile aggregate search cost.

Figure 6: Search Cost Distributions




startup search costs. Panel A shows that, despite being simple, our model with startup costs is able to recreate the amplitude and frequency of search cycles, with the amplitude being somewhat underestimated. Importantly, the full model precisely fits the sharp and permanent shift in search intensity after the temporary shock to search incentives caused by the price war.

Panel B of Figure 7 yields a stark contrast for the restricted model without startup costs. In this model, the distribution of recurrent search costs is identified by variation in search activity with search benefits, both when those benefits are unprecedented and when they are moderate. While the model is able to capture the amplitude and frequency of search cycles, the model is unable to produce a permanent shift in search behavior after the shock. As the gains from search return to baseline following the shock, in the absence of startup costs, the standard non-sequential search model predicts search intensity will also return to baseline.

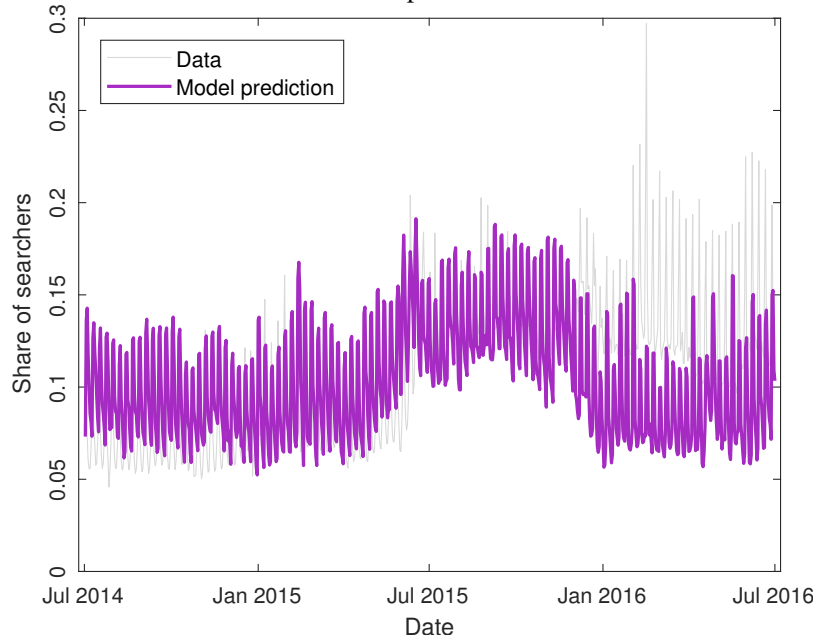
Finally, panel C shows the predicted evolution of consumer experience with search. In grey (left scale), we depict the population-weighted average gains from search, and in the foreground (right scale), we show the predicted fraction of active consumers. The fraction of consumers who have incurred startup search costs is approximately a step function over time. When the gains from search are abnormally high in the middle of the sample, the predicted fraction of consumers with search experience grows rapidly from 12% to 18%.

Figure 7: Price and Search Cycles

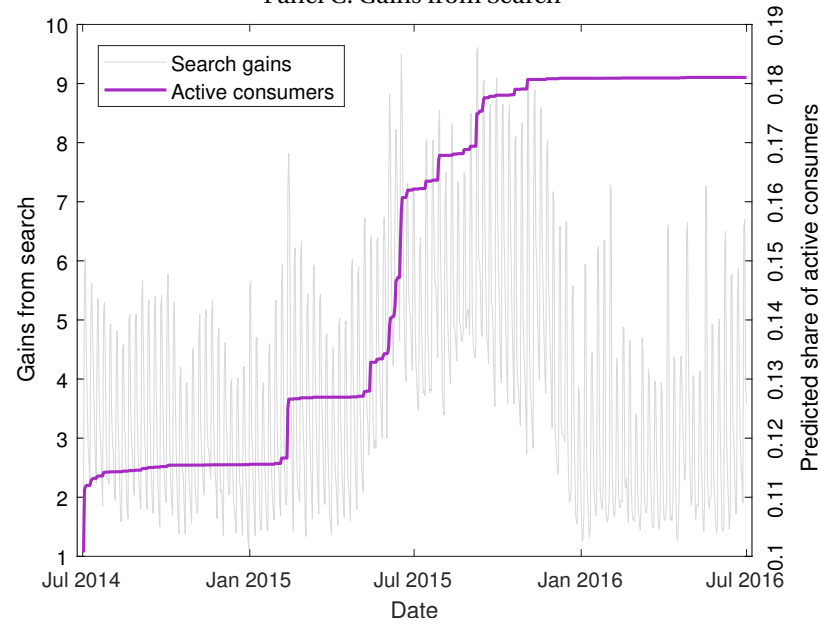
Panel A: Non-Sequential Search Model
with Startup Search Costs



Panel B: Non-Sequential Search Model
without Startup Search Costs



Panel C: Gains from Search



Panel C further helps to explain the differential performance of the two models. In the full model, following the price war, a greater fraction of consumers have experience with search. This changes the aggregate relationship between search activity and search benefits, and explains why search remains elevated after search benefits decline. By contrast, the restricted model does not allow for a dynamic relationship between search activity and search benefits, and therefore predicts that search activity must rise and fall with current search gains.

5 Summary and discussion

In studying search frictions, economists have, to date, assumed that search costs are independent of search history. Exploiting a natural experiment in retail gasoline, together with unique data on retail prices and search behavior, we have provided evidence on a new form of search costs that we call startup search costs. The novel evidence of history dependence in search behavior that we find suggests that the first search cost sunk by a consumer is drastically different from subsequent search costs. Search experience matters.

The results from our simple empirical search model highlight the implications of startup search costs for the measurement of search frictions. Model misspecification that ignores startup search costs yields overestimated recurrent search costs. Moreover, we have shown how a standard non-sequential search model is unable to account for a permanent rise in market-level search from a temporary exogenous shock to search incentives.

Because we work with market-level and not individual-level search data, we are unable to identify deeper microfoundations for startup search costs. We can think of at least four possible mechanisms to explore in future research. First, startup search costs could be driven by technology adoption costs (Foster and Rosenzweig, 2010) with online price clearinghouses. Second, startup search costs could reflect consumers holding biased beliefs about the value of search (Koulayev and Alexandrov, 2017), and updating these beliefs after trialing the clearinghouse during the price war. Third, consumers could rapidly form habits after trialing the clearinghouse (Becker and Murphy, 1988) with minimal rates of habit decay. Finally, the inertia in non-adoption of the clearinghouse for 15 years before the war, followed by a permanent shift in usage after, could reflect procrastination or time-inconsistency (O'Donoghue and Rabin, 1999) in trialing and learning to use the clearinghouse.

Understanding the role of startup search costs and their underlying mechanisms is important for policy. We believe that the evidence presented here points to a new and important policy challenge with online search platforms aimed at promoting price transparency and market efficiency: policymakers need to get consumers “over the hump” in starting to use such platforms. We have found that this hump prevented consumers from engaging with a well-

established price clearinghouse for 15 years. It took a three-week, temporary price shock to substantially increase online price search by 70% more than a decade and a half after the clearinghouse was established. The lesson for policy is that large, temporary shocks to search incentives can help consumers overcome startup search costs and lead to long-run adoption of search platforms. Policy interventions that encourage customers to experiment with such platforms are a potential remedy for overcoming startup search costs.

On the supply-side, our study raises a separate question for future research: what is the impact of startup search costs for firms' pricing decision? In the retail gasoline market that we study, we obtain an interesting implication for firms' pricing decisions. As we have mentioned, in [Byrne and de Roos \(2019\)](#) we document evidence consistent with tacit collusion in this market. In this context, by encouraging consumers to engage with search, the temporary price war may have led to an increase in demand elasticity, and therefore collusive outcomes may have become more difficult to sustain following the war. This suggests a new trade-off – price variation could lead to a sustained increase in demand elasticity – facing cartel members contemplating either an adjustment of pricing policies or defection from a cartel.

Appendix

A Consumers with a dynamic perspective

In the model of Section 4, we presumed consumers take a static perspective when deciding whether to incur startup search costs. In this section, we illustrate the implications of relaxing this assumption. Consider the perspective of consumer i who evaluates the impact of today's search decision on the search environment that she will face in the future. To fix ideas, we begin by laying out the Bellman equation faced by consumer i at time t when she adopts this dynamic perspective.

Given current search state w_{it} , her type τ_i , and price vector \mathbf{p}_{it} , her current valuation is given by

$$V(w_{it}) = \max_{\chi_{it} \in \{0,1\}} \chi_{it} (\bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - s_{it}) + (1 - \chi_{it}) (\bar{u} - \kappa_i \bar{p}_{it}) + \delta \mathbb{E}_t V(w_{it+1}),$$

$$w_{it+1} = w_{it} + (1 - w_{it}) \chi_{it},$$

where $\chi_{it} = 1$ indicates a decision to search today and $\chi_{it} = 0$ indicates no search; and \mathbb{E}_t indicates period- t expectations over future price distributions.²² The parameter δ describes the rate at which consumers discount the next fuel purchase, and could reflect impatience, and concerns about the decay or obsolescence of current knowledge of the search process.

Expectations of future prices play an important role through their influence on the continuation value of the consumer's dynamic problem. For illustration, we consider a simple expectations process. We say that consumer i adopts stationary expectations if she anticipates the current price distribution to be observed in subsequent periods: $\mathbf{p}_{t+k} = \mathbf{p}_t$, for $k > 0$. This leads to the following proposition.

Proposition 1. *Suppose consumer j adopts a static perspective with search costs c_j and f_j , and consumer i adopts a dynamic perspective with stationary expectations and search costs c_i and f_i . Then consumers i and j are observationally equivalent if $c_i = c_j$ and $f_i = f_j / (1 - \delta)$.*

Proof. First, consider consumer j . Based on her static perspective, she searches iff $g_{jt} > s_{jt}$. If $w_{jt} = 1$, she searches iff $g_{jt} > c_j$; if $w_{jt} = 0$, she searches iff $g_{jt} > c_j + f_j$.

Next, consider consumer i and suppose $w_{it} = 1$. In this case, she has already sunk her startup search costs and, as a result, her current search decision has no dynamic consequences.

²²For simplicity, we ignore the intertemporal search opportunities presented by the Fuelwatch program in this formulation, and consider search gains in period t based solely on the period t price distribution.

Thus, she chooses to search iff $g_{it} > c_i$. Because $w_{it} = 1$ is an absorbing state, we can solve for the value $V(1) = u(c_i)/(1 - \delta)$, where $u(\cdot)$ is defined in equation (1).

Suppose instead $w_{it} = 0$ and observe that consumer i has value

$$V(0) = \max_{\chi_{it} \in \{0,1\}} \chi_{it} (\bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - c_i - f_i) + (1 - \chi_{it}) (\bar{u} - \kappa_i \bar{p}_{it}) \\ + \delta (\chi_{it} V(1) + (1 - \chi_{it}) V(0)).$$

Consumer i searches in period t iff

$$\bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - c_i - f_i + \delta \frac{u(c_i)}{1 - \delta} > \bar{u} - \kappa_i \bar{p}_{it} + \delta V(0).$$

Observing that consumer i makes the same decision whenever $w_{it} = 0$, we can deduce that she searches iff

$$(1 - \delta) (\bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - c_i - f_i) + \delta u(c_i) > \bar{u} - \kappa_i \bar{p}_{it}. \quad (3)$$

Next, we show that, when $w_{it} = 0$, consumer i searches iff $g_{it} > c_i + (1 - \delta)f_i$. We break this into two steps. First, observe that if $g_{it} > c_i$, then a consumer who had already sunk her startup search costs would choose to search. This means that $u(c_i) = \bar{u} - \kappa_i \min\{\mathbf{p}_{it}\} - c_i$. Substituting into (3) leads to the conclusion that i searches iff $g_{it} > c_i + (1 - \delta)f_i$. Next, suppose that $g_{it} \leq c_i$. In this case, $u(c_i) = \bar{u} - \kappa_i \bar{p}_t$. Suppose further that $\chi_{it} = 1$. Substituting into (3) leads to the condition $g_{it} > c_i + f_i$, a contradiction. Therefore $\chi_{it} = 0$ whenever $g_{it} \leq c_i$. Combining the two cases, we have our desired result that consumer i searches iff $g_{it} > c_i + (1 - \delta)f_i$.

Finally, comparing consumers i and j leads to the conclusion that their choices are identical if $c_i = c_j$ and $(1 - \delta)f_i = f_j$, as required. \square

Proposition 1 gives a feeling for the impact of the consumer's perspective on inferences about search costs under the assumption of stationary expectations. The perspective adopted by consumer i has no impact on the inferences we make about her recurrent search costs. However, particularly for patient consumers, inferred startup search costs will be substantially higher if we presume consumers adopt a dynamic perspective.

The logic of the proof of Proposition 1 provides an indication of the impact of the assumption of stationary expectations. Suppose that in period t , consumer i decides to first engage in search. Under the stationarity assumption, she anticipates that she would also have chosen to initiate search in period $t + 1$ had she not chosen to search in period t . Thus, she derives a benefit of f_i in every subsequent period. Similarly, if instead she anticipates that price variation and the gains to search will increase over time, then she will also anticipate engaging in search

in each period, and the value to her of initiating search will be the same. Alternatively, if she expects the gains to search to fall, she may anticipate that there are future periods in which she would not be willing to initiate search. In this case, by assuming stationarity, startup search costs will be overestimated.

B Specification of search gains: robustness

Recall that the gains to search for consumer i at time t are defined by

$$g_{it} = (\bar{p}_{it} - \min\{\mathbf{p}_{it}\}) \times \kappa_i.$$

In the body of the paper, we presumed that average and minimum prices for consumer i were taken with respect stations within a 5km radius of the centroid of her local district. In this section, we also consider search radii of 2km and 10km.

As the search radius varies between 2km and 10km, there is a noticeable impact on the consumer's choice set. With a search radius of 2km, 5km, and 10km, there are on average 1.96 stations (s.d. 1.57), 10.29 stations (s.d. 6.32), and 34.26 stations (s.d. 18.86) within the choice radius for each district. For estimation purposes, we eliminate all districts with less than two stations inside the search radius.

For comparability, we retain the same format for the presentation of results. Table 3 contains estimation results. Columns (1) and (2) ((3) and (4)) contain estimates based on a 2km (10km) search radius. For each specification of the search radius, the left and right columns contain, respectively, estimates based on the full model and a constrained model that does not include startup search costs. Figures 8 and 9 depict, for the 2km search radius specification, the cumulative distribution of estimated startup and recurrent search costs, and predictions for the model, respectively. Figures 10 and 11 contain analogous information for the 10km search radius specification.

The main qualitative features we highlighted earlier carry over to alternative specifications of the search radius. In particular, the model with startup costs leads to qualitatively different inferences over recurrent search costs, the fit of the model is much improved by incorporating startup search costs, and only the model with startup search costs is able to explain the permanent increase in search activity following the temporary shock to search gains.

Adjusting the search radius does lead to some variation in results. When the search radius is reduced, measured search gains tend to be lower. This can be seen by comparing Panel C of Figures 7, 9, and 11. As a result, to rationalize observed search activity, estimated startup and recurrent search costs are lower when the search radius is reduced. This is best seen by

Table 3: Estimation Results with Alternative Search Radii

	2km Search Radius		10km Search Radius	
	With Startup Costs (1)	Without Startup Costs (2)	With Startup Costs (3)	Without Startup Costs (4)
Recurrent search cost distribution				
μ_c	0.205 (0.017)	0.455 (0.005)	0.215 (0.017)	0.421 (0.003)
σ_c	16.907 (2.341)	176.596 (6.759)	36.781 (3.829)	1322.721 (17.370)
Startup search cost distribution				
μ_f	1.945 (0.081)		6.201 (0.144)	
σ_f	12.110 (0.613)		3.811 (0.135)	
Objective function, $G(\hat{\theta})$	0.374	0.816	0.375	0.934

Notes: Robust standard errors are in parentheses (). The number of observations is $T = 731$ dates. All calculations assume consumers purchase 50 liters of gasoline.

comparing Figures 6, 8, and 10. Consider first the model with startup search costs. Estimated 20th percentile startup costs are \$9.54, \$13.30, and \$15.51 when the search radius is 2km, 5km, and 10km, respectively. At the left tail of the distribution, estimated recurrent search costs are negligible in each specification of the search radius. The change to our search gain specification makes a bigger difference to estimates for the model without startup search costs. The 20th percentile consumer now has estimated recurrent search costs of \$3.98, \$9.94, and \$22.06 when the search radius is 2km, 5km, and 10km, respectively.

Comparing Figure 9 and 11 reveals some subtle differences in model predictions. When search gains are defined more locally, this accentuates the volatility in search gains, and there is an associated increase in the high-frequency volatility in predicted search, both for the full model (Panel A) and the restricted model without startup search costs (Panel B). Finally, Panel C also suggests that, when search gains are defined more locally, to rationalize the volume of search activity, a greater proportion of consumers are predicted to incur their startup search costs.

Figure 8: Search Cost Distributions, 2km Search Radius

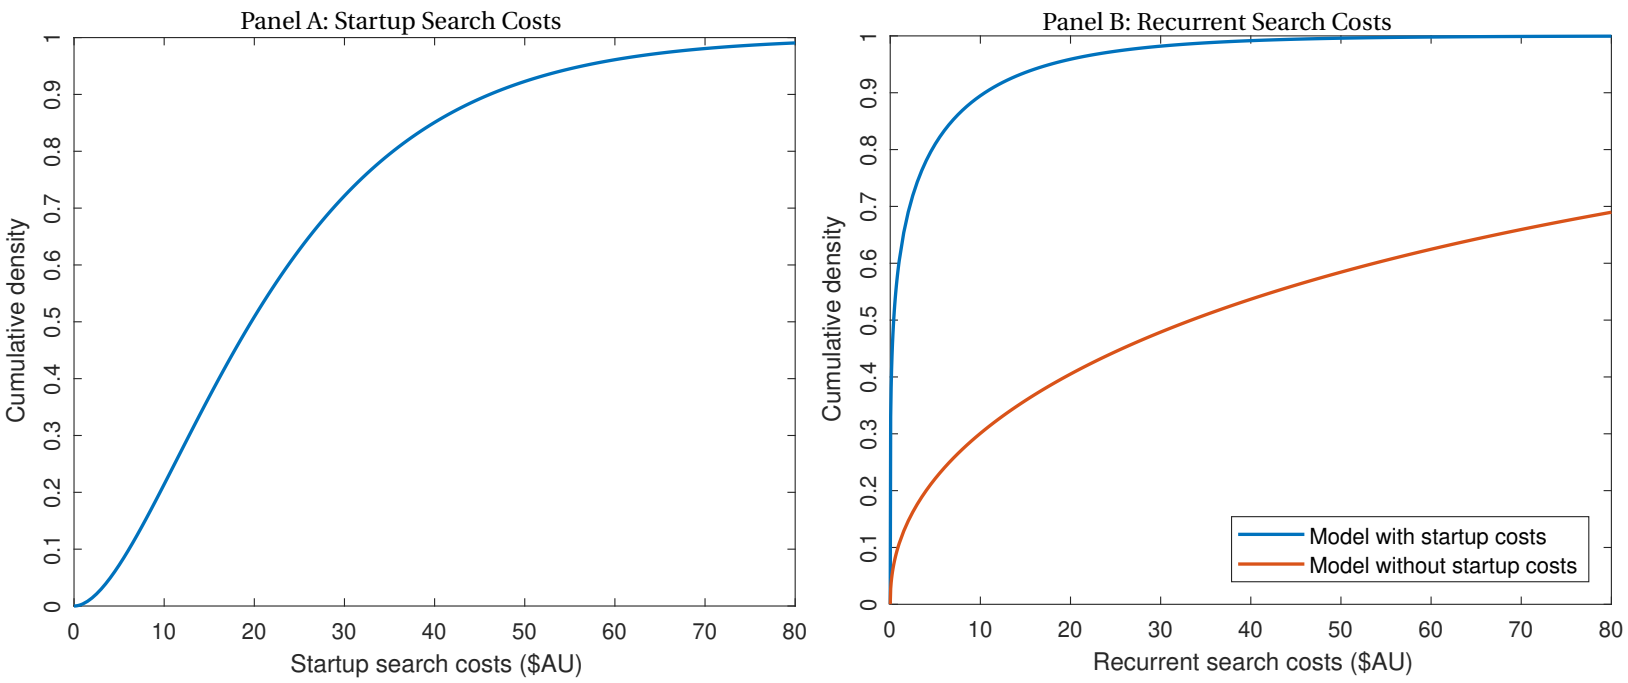

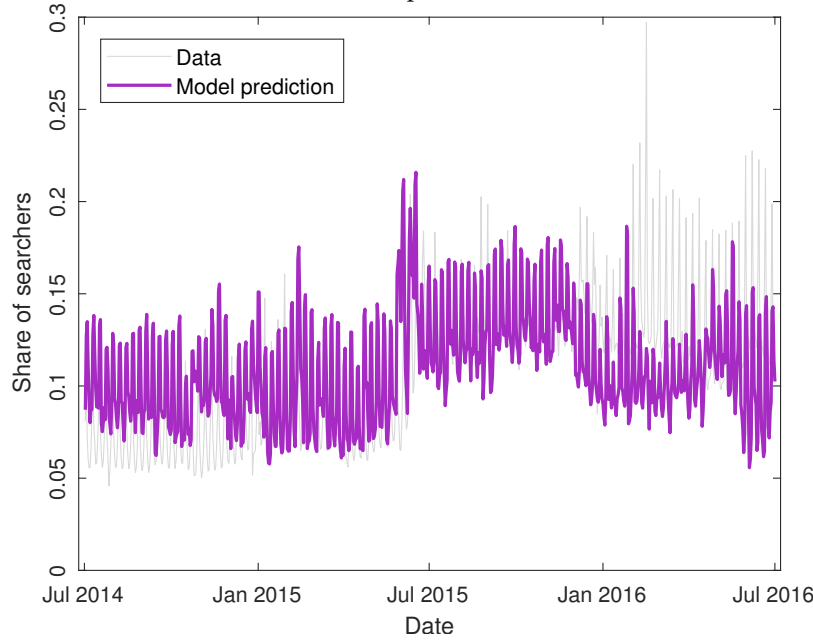


Figure 9: Search Model Predictions, 2km Search Radius

Panel A: Non-Sequential Search Model
with Startup Search Costs



Panel B: Non-Sequential Search
without Startup Search Costs



Panel C: Gains from Search

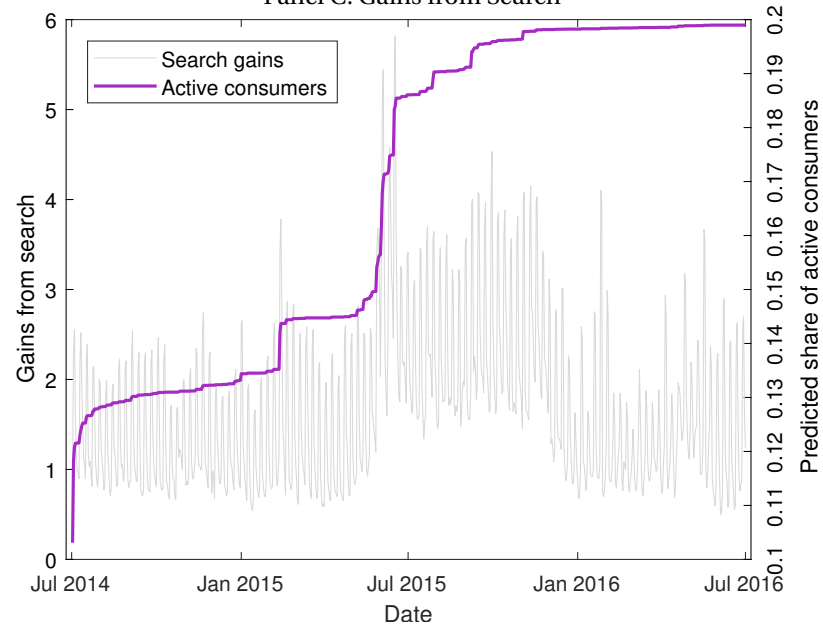


Figure 10: Search Cost Distributions, 10km Search Radius



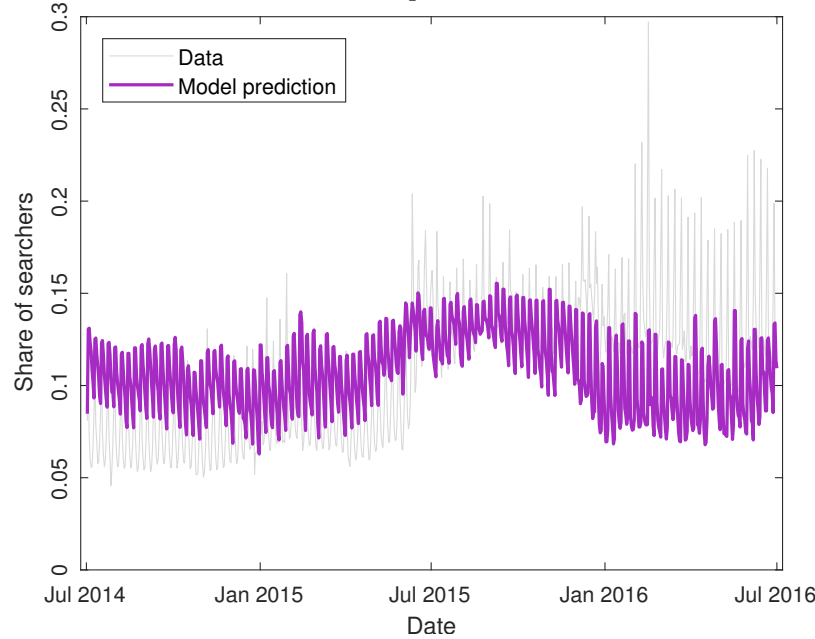


Figure 11: Search Model Predictions, 10km Search Radius

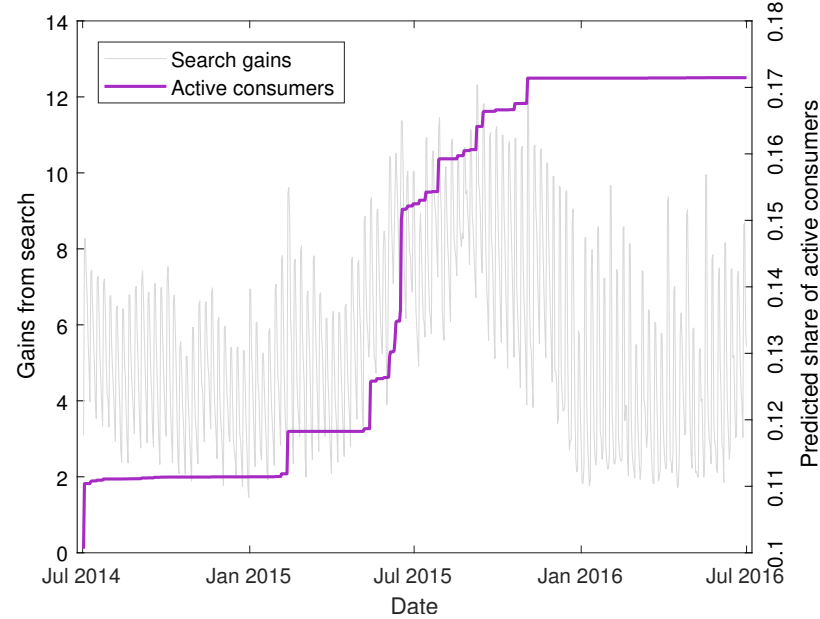
Panel A: Non-Sequential Search Model
with Startup Search Costs



Panel B: Non-Sequential Search
without Startup Search Costs



Panel C: Gains from Search



References

- Australian Competition and Consumer Commission. Petrol prices and Australian consumers. Report of the ACCC Inquiry into the Price of Unleaded Petrol, 2007.
- Australian Competition and Consumer Commission. Petrol prices and Australian consumers. Report of the ACCC Inquiry into the Price of Unleaded Petrol, 2014.
- Australian Competition and Consumer Commission. Report on the Australian petroleum market, September quarter 2017, 2017.
- Michael R. Baye, John Morgan, and Patrick Scholten. Information, search, and price dispersion. 2006. Chapter 6 in *Handbook in Economics and Information Systems Vol. 1* (T. Hendershott, Ed.), Amsterdam: Elsevier.
- Gary S. Becker and Kevin M. Murphy. A theory of rational addiction. *Journal of Political Economy*, 96:675–700, 1988.
- Tom Blake, Dominic Coey, Kane Sweeney, and Steve Tadelis. Price salience and product choice. NBER Working paper 25186, 2018.
- Nicholas Buchholz. Spatial equilibrium, search frictions and dynamic efficiency in the taxi industry. mimeo, Princeton University, 2018.
- David P. Byrne and Nicolas de Roos. Consumer search in retail gasoline markets. *Journal of Industrial Economics*, 65:183–193, 2017.
- David P. Byrne and Nicolas de Roos. Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109:591–619, 2019.
- Babur de Los Santos, Ali Hortascu, and Matthijs R. Wildenbeest. Testing models of consumers search using data on web browsing and purchasing behavior. *American Economic Review*, 102:2955–2980, 2012.
- Andrew Eckert. Empirical studies of gasoline retailing: A guide to the literature. *Journal of Economic Surveys*, 27:140–166, 2013.
- Sara Fisher Ellison. Price search and obfuscation: an overview of the theory and empirics. *Handbook of the Economics of Retail and Distribution*, 2016. Emek Basker (eds), Elgar.
- Andrew D. Foster and Mark R. Rosenzweig. Microeconomics of technology adoption. *Annual Review of Economics*, 2:395–424, 2010.

- Han Hong and Matthew Shum. Using price distributions to estimate search costs. *RAND Journal of Economics*, 37, 2006.
- Sergei Koulayev. Search for differentiated products: Identification and estimation. *Rand Journal of Economics*, 45:553–575, 2014.
- Sergei Koulayev and Alexei Alexandrov. No shopping in the u.s. mortgage market: Direct and strategic effects of providing information. CFPB Working Paper 2017-01, 2017.
- Laurence Levin, Matthew S. Lewis, and Frank A. Wolak. High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9, 2017.
- Ted O’Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89: 104–124, 1999.
- Dieter Pennerstorfer, Philipp Schmidt-Dengler, Nicolas Schutz, Christoph Weiss, and Biliana Yontcheva. Information and price dispersion: Theory and evidence. CEPR Discussion Paper 10771, 2018.
- Stephan Seiler and Fabio Pinna. Estimating search benefits from path-tracking data: Measurements and determinants. *Marketing Science*, 36:565–589, 2017.
- Meredith Startz. The value of face-to-face: Search and contracting problems in nigerian trade. mimeo, Stanford University, 2018.
- George J. Stigler. The economics of information. *Journal of Political Economy*, 69:213–225, 1961.
- Matthijs R. Wildenbeest. An empirical model of search with vertically differentiated products. *Rand Journal of Economics*, 42:729–757, 2011.