Online Appendix

Startup Search Costs

David P. Byrne and Nicolas de Roos

October 9, 2020

Contents

A Testing for breaks in prices, margins, and search

In this Appendix, we formally test for structural breaks in price levels, margin levels, price dispersion, and search levels in the time series presented in Panels A-D of Figure 4 in the paper.

A.1 Graphs with standard errors

In Figure A.1 on the next page, we reproduce Panels A-D of Figure 4 except we plot the weekly sample average of a given price, margin or search variable, plus or minus one standard error of the weekly sample average of that variable.[1] The scales of the axes are identical to those in Figure 4 to facilitate comparison. These plots yield similar insights to those from Figure 4. Panels A and B show weekly price levels and margins do not change substantially after the price war. In contrast, panels C and D reveal a temporary spike in price dispersion during the price war with a concurrent large and permanent shift in search intensity.

A.2 Testing for difference in means pre/post price war

We conduct two-sample t-tests of differences in means between the pre-war period (July 1, 2014-May 27, 2015) and the post-war period (July 16, 2015-June 30, 2016).[2] We also conduct two-sample t-tests based on a longer-run post-price war period of January 1, 2016-June 30, 2016 to test for long-run shifts in means.

Table A.1 presents our test results, where we use Newey and West (1987) standard errors with 14 lags to account for persistence in the variables. Consistent with the patterns from Figure 4, we find a -12.09 cpl and -14.22 cpl difference in means in average retail and wholesale prices in the post-war period, both of which are non-negligible in magnitude. Because prices do not fall as fast as wholesale prices over time, we find a 2.13 cpl increase in daily average margins in the post-war relative to the pre-war period that is also statistically significant. Notably, the bottom panel of the table shows that margin levels exhibits a statistically insignificant difference of 1.19 cpl relative to the pre-war mean over a longer time horizon.

Turning to price dispersion in Table A.1, we find a statistically significant 1.24 cpl increase in the mean of the daily standard deviation of prices. However, in the bottom panel of the table

[1] To compute the standard error for weekly average prices and margins we first compute daily averages for these variables; this is what is plotted in grey scale in Figure 4. The standard error for the weekly mean price/margin is then computed using the seven daily averages for prices/margins in a given week. Likewise, the standard error for average daily price dispersion in a given week is computed based on the daily price dispersion for the seven days in that week. Finally, the standard error for weekly average search is computed based on the seven daily website visits for a given week.

[2] Here, we exclude the price war window in conducting our tests of the difference of means. Our structural break tests in Section A.3 below will test for an arbitrary break point in the data, allowing for the break in a sample mean to occur anywhere before, within, or after the price war period.

Figure A.1: Prices, Margins and Search Before and After the Price War

Plotting Weekly Sample Average Plus or Minus 1 Standard Error of the Weekly Sample Average

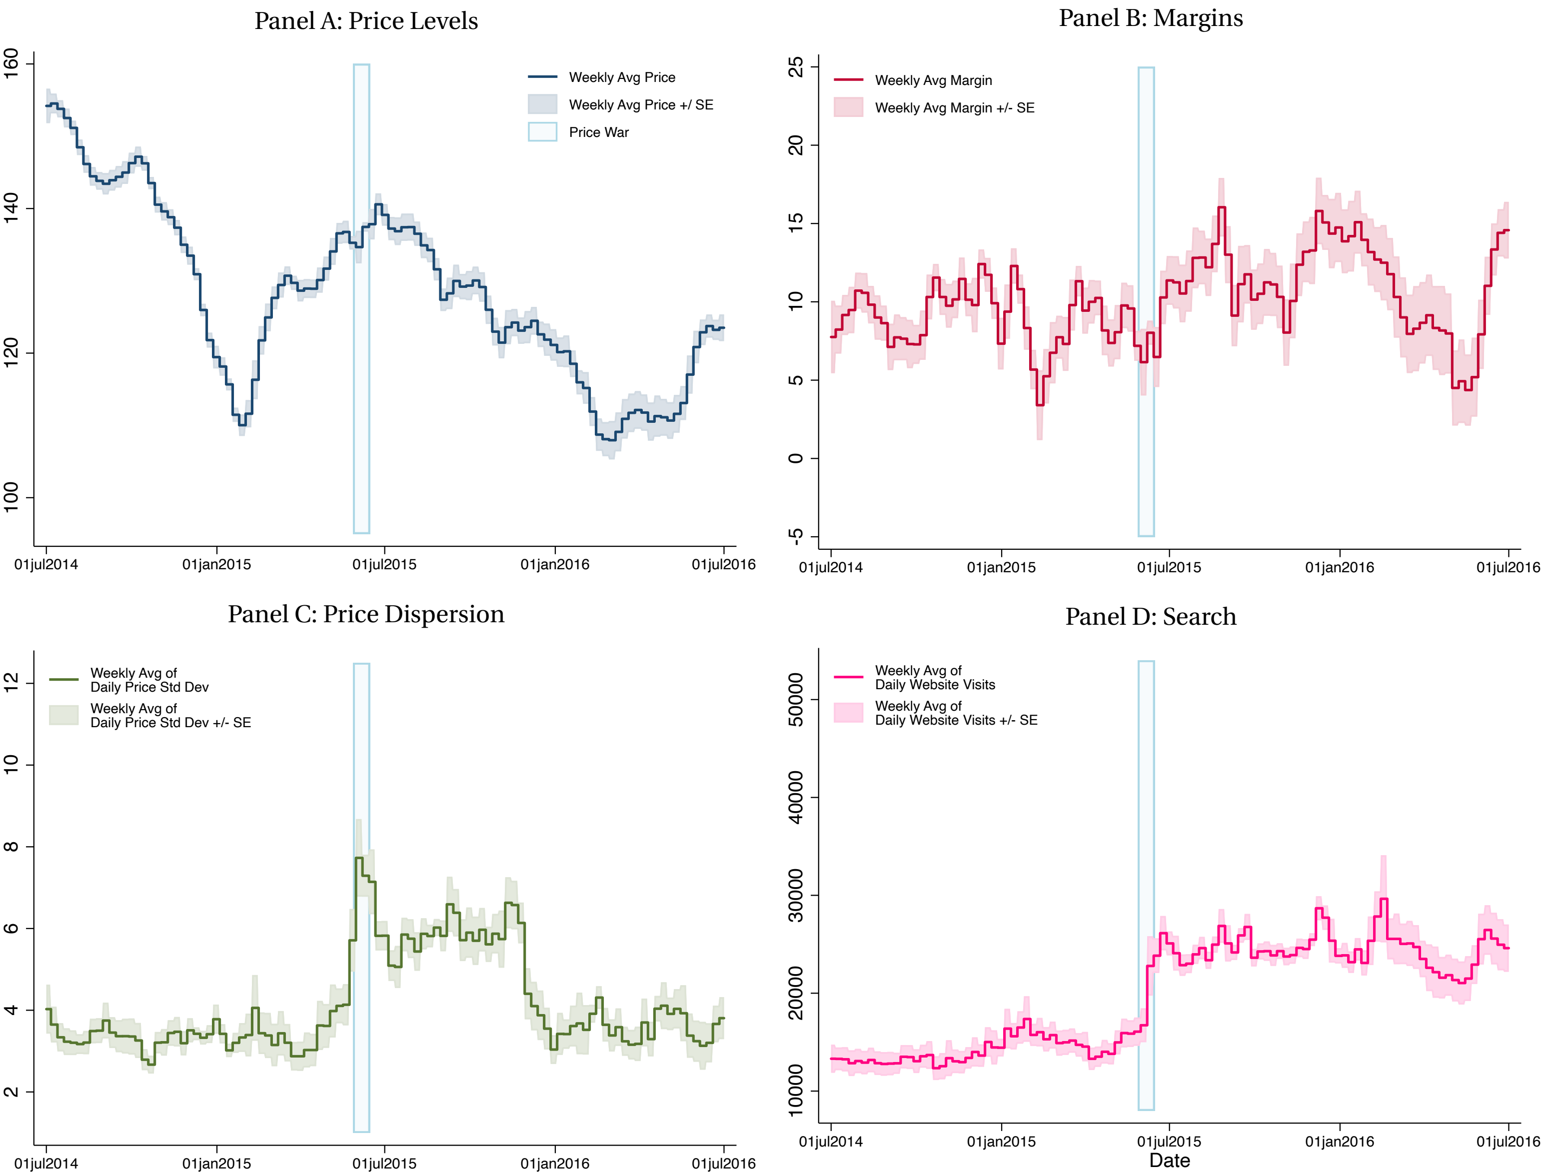

where we consider a longer-run post-war window, we find a statistically insignificant 0.19 cpl increase in the daily standard deviation of prices. The difference in results is driven by the fact that while daily price dispersion falls immediately after the price war, it remains higher than baseline (pre-war) levels until the end of 2015. Visually, this can be seen in panel (C) of Figure 4 in the paper.[3] Statistically, theses findings confirm that the level of daily price dispersion eventually returns to its baseline (pre-war) level over our sample period.

[3]We explain why this happens in Appendix B below.

Table A.1: Two-Sample t-Tests for Differences in Means Before and After the Price War

Daily Variable	Pre-War Mean	Post-War Mean	Diff. in Means	p-value
Full-Sample Pre/Post War Differences				
Avg. Price	134.67	122.58	-12.09	0.000
	(0.68)	(0.55)	(2.95)	
Avg. Wholesale Price	125.67	111.45	-14.22	0.000
	(0.65)	(0.46)	(2.86)	
Avg. Margin	9.00	11.12	2.13	0.001
	(0.22)	(0.31)	(0.63)	
Price Std. Dev.	3.39	4.63	1.24	0.000
	(0.06)	(0.09)	(0.23)	
Website Visits	14170	24541	10371	0.000
	(194)	(245)	(399)	
Longer-Run Pre/Post War Differences				
Avg. Price	134.67	114.64	-20.02	0.000
	(0.68)	(0.57)	(2.74)	
Avg. Wholesale Price	125.67	104.45	-21.22	0.000
	(0.65)	(0.34)	(2.66)	
Avg. Margin	9.00	10.19	1.19	0.208
	(0.22)	(0.49)	(0.95)	
Price Std. Dev.	3.39	3.58	0.19	0.076
	(0.06)	(0.09)	(0.11)	
Website Visits	14170	24189	10019	0.000
	(194)	(447)	(591)	

Notes: In the *Full Sample* panel, the pre-war period is July 1, 2014-May 27, 2015 and the post-war period is July 16, 2015-June 30, 2016. In the *Longer-Run* panel, the pre-war period is July 1, 2014-May 27, 2015 and the post-war period is January 1, 2016-June 30, 2016. Newey and West (1987) standard errors with 14 lags presented in parentheses for difference in means, with corresponding p-value for the two-sided test of the null of equality of pre-war and post-war means presented in the last column.

Finally, Table A.1 reports a significant and large break of 10,371 additional websites visits per day after the price war, which is a 73% increase in the mean of daily search relative to the pre-war mean. Unlike the change in price-dispersion, however, we find this difference is stable irrespective of the sample window used for the post-war period, as shown by the test results in the bottom panel of the table. Over the longer time horizon in the bottom panel we find a statistically significant and large 10,019 increase in the mean of daily website visits.

A.3 Structural break tests

To finish our analysis of pre/post war changes in prices, margins, and search, we conduct Andrews (1993) structural break tests with each of these respective time series. These tests allow for a structural break in the mean of a variable over time that does not *a priori* specify at what date a break occurs, if one exists.

To implement the test, we estimate regressions of the following form:

$$y_t = \beta_0 + \beta_1^\tau 1\{t \geq \tau\} + \epsilon_t \tag{1}$$

where y_t is an outcome variable on date t, $1\{t \geq \tau\}$ is a dummy variable equaling one for dates on and after date τ, and ϵ_t is the econometric error. The coefficient β_1^τ corresponds to a level break in y_t after τ, and we can test for a break given τ based on the null $H_0: \beta_1^\tau = 0$ vs. $H_1: \beta_1^\tau \neq 0$. We implement the Andrews (1993) supF test by estimating (1) for a range of τ values over our sample period, testing H_0 vs H_1 for each candidate τ value, and obtaining the F-statistic for each test, F_τ. The supF test statistic is the supremum of F_τ, and the τ that corresponds to the supF test statistic identifies the unknown structural break in the data.

To focus on long-run level breaks in our pricing, margin, and search variables, as well as to simplify computation, we conduct our supF tests based on weekly data. For price and margin levels, this means conducting the tests based on weekly average prices and margins. For price dispersion, the means first computing daily price dispersion and then computing the weekly average of this variable to obtain a weekly average of daily dispersion with which to conduct the test. Finally, for search this means computing the weekly average of daily visits to the FuelWatch platform and conducting the test based on this average.

In Panels A-H of Figure A.2 below, we presents our supF test results for price levels, margin levels, price dispersion, and search. For each variable we produce two graphs. The first plots the F-statistic F_τ and p-value for the test $H_0: \beta_1^\tau = 0$ vs. $H_1: \beta_1^\tau \neq 0$ for all possible break points τ. We highlight the location of the supF and hence the estimated timing of the unknown break in the data. In the second graph, we plot the corresponding break point estimates $\hat{\beta}_1^\tau$ along with their 95% confidence intervals for all possible τ values. We again highlight the location of the break point as determined by the supF test.

Starting with price levels, panels (A) and (B) identifies the last week of October 2014 as the break point, with a break estimate of -22.85 cpl. The supF test identifies a break point early in the sample period because of the downward trend in price levels, which can be seen in panel A of Figure 4. As the figure shows, this is driven by the downward trend in wholesale oil prices over the sample period.

In panels C and D of Figure A.2 we find a break point in margins in the third week of June

2015, just after the price war. The estimated break in margins is 2.31 cpl which is statistically significantly different from 0 with $p = 0.019$. However, the F-statistic and p-value plots in panel C as well as the estimated breaks in panel (C) show there does not exist a break in margins over longer time horizons in 2016. This was foreshadowed by our longer-time horizon two-sample test results for a break in margins from Table A.1 above.

Panels E and F of Figure A.2 presents a noisy picture of where a break occurs for daily price dispersion. The structural break is identified in the third week of April 2015, well-before the price war. The reason why the supF identifies this week as the break point is because it is able to group the entire sudden large jump in price dispersion and subsequent gradual decline in dispersion within an entire sub-sample after April 2015. This creates the largest possible break in daily price dispersion levels, and the supF statistic zeros in on this point.

The decline in the F-statistic and rise in the p-value for the supF test in panel E of Figure A.2 further shows how break points in price dispersion within the price war period do not yield statistically significant breaks in price dispersion. This is because these candidate break points split the price war period into the first and second half of the sample period. By averaging over high and low price dispersion periods between the baseline and price war period, the associated break in price dispersion shrinks, causing F_τ to fall for τ values within the price war period.

The upshot of this discusssion is that the supF statistic is not as useful for identifying a structural break in price dispersion in our sample precisely because a permanent structural break *does not exist*. That price dispersion exhibits a temporary shock is visually clear in panel C in Figure 4. This is confirmed in our two-sample t-test in the bottom panel of Table A.1 above where we fail to find a break in daily price dispersion levels over longer time horizons with dispersion eventually returning to baseline (pre-war) levels between January 1 and July 1, 2016.

Finally, panels G and H of Figure A.2 present the structural break results for search. A clear picture emerges as the supF test zeroes in on the structural break in search that occurs within the price war period (panel G) with a precisely-estimate break in the daily mean of search of 10,273 website visits on the FuelWatch platform (panel H).

Figure A.2: Structural Break Test Results

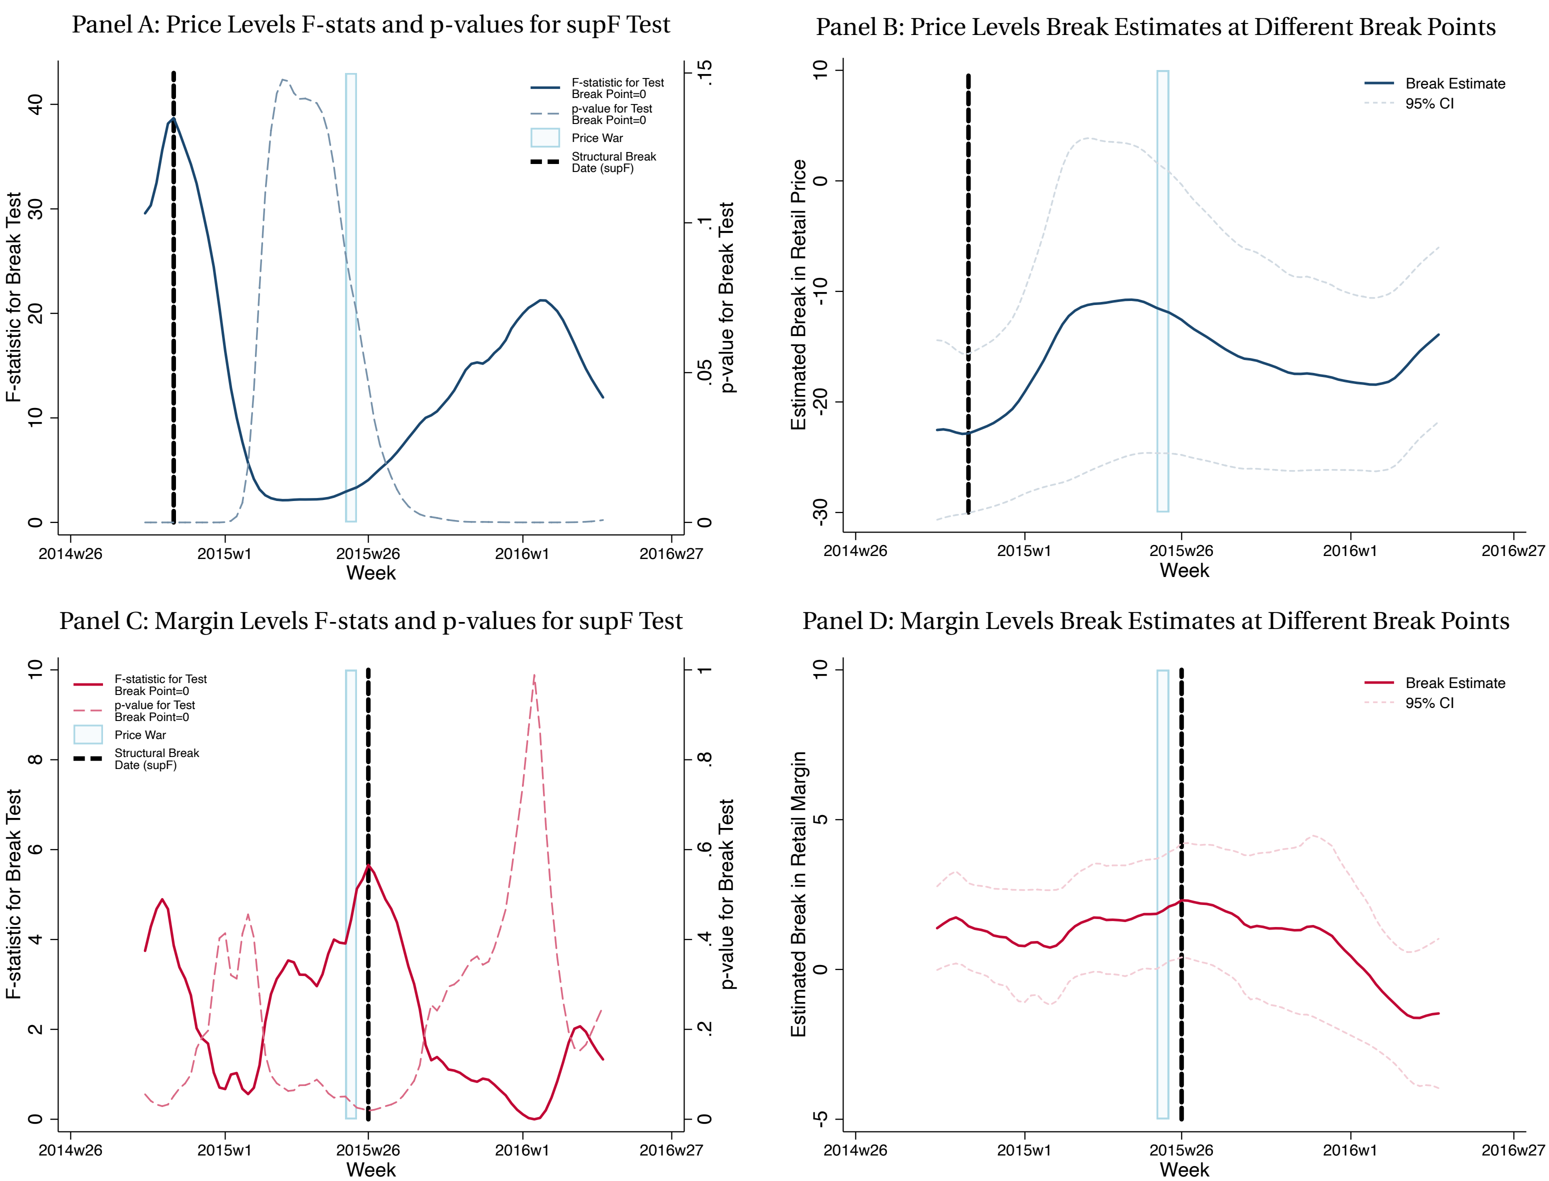

Structural Break Test Results (continued)

Panel E: Daily Price Std. Dev. F-stats and p-values for supF Test

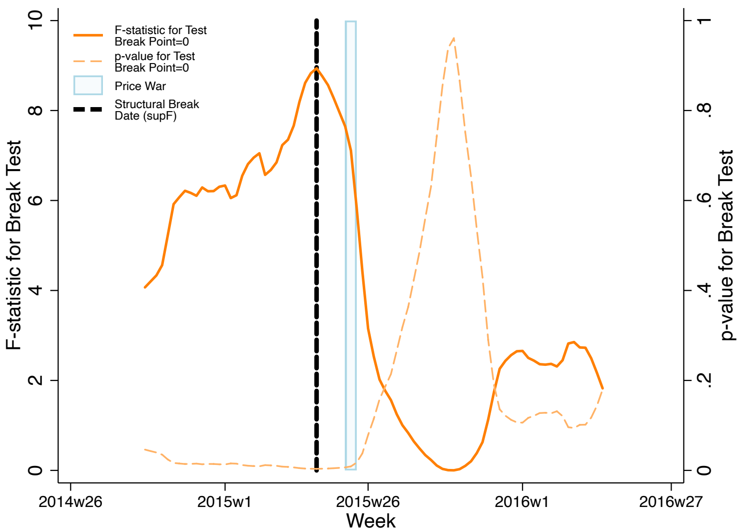

Panel F: Daily Price Std. Dev. Break Estimates at Different Break Points

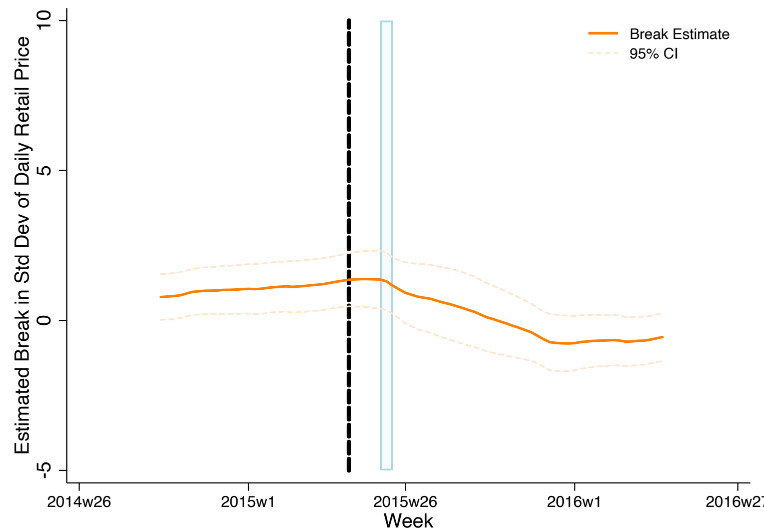

Panel G: Search F-stats and p-values for supF Test

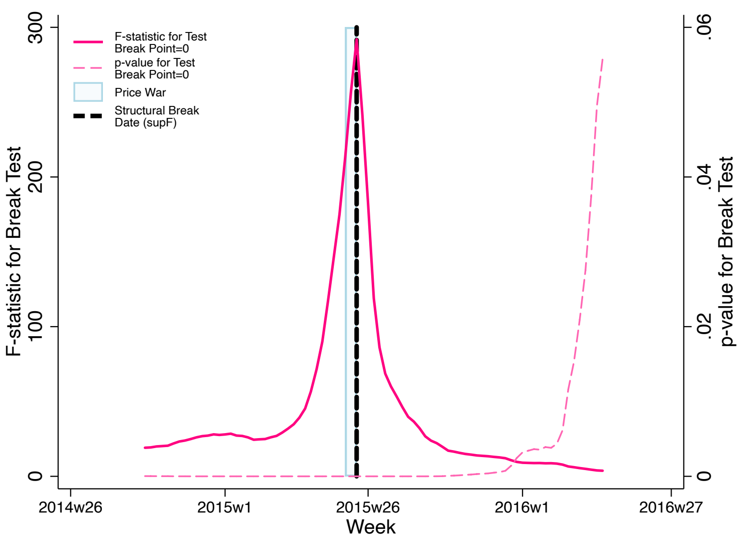

Panel H: Search Break Estimates at Different Break Points

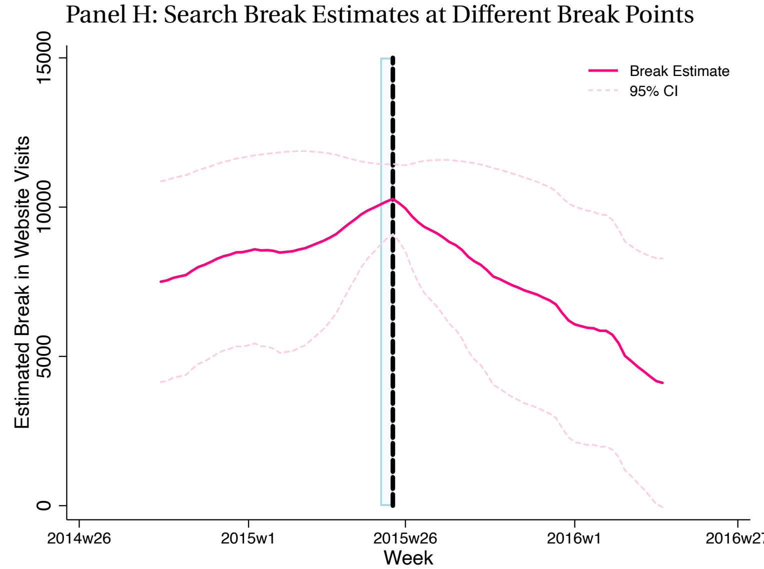

Unique searchers and search intensity

We further conduct supF tests for structural breaks for the number of unique visitors to FuelWatch per month and number of visits per unique visitor to FuelWatch. Recall from Figure 5 in the paper that we document a jump in the former and a lack of a jump in the latter following the price war. In Panels A-D of Figure A.3, we present supF structural break tests for breaks in these variables. Panel A shows that the break in the number of unique visitors corresponds precisely to the price war in June 2015. Indeed, panel B reveals a statistically significant 91,333 increase in the number of unique visitors to the FuelWatch platform per month at the onset of the price war. This represents a substantial 61% increase in unique visitors relative to a pre-war mean of 150,620 unique visitors per month.

Panels C and D of Figure A.3 examine breaks in the number of visits per unique visitor per month. Here, we identify a structural break in July 2015, just after the price war. The size of the break is 0.29 additional visits per month, which is also statistically significant. It is, however, a relatively small increase compared to its pre-war sample mean of 2.85 visits per unique visitor per month.

Figure A.3: Structural Break Test Results for Unique Visitors and Search Intensity

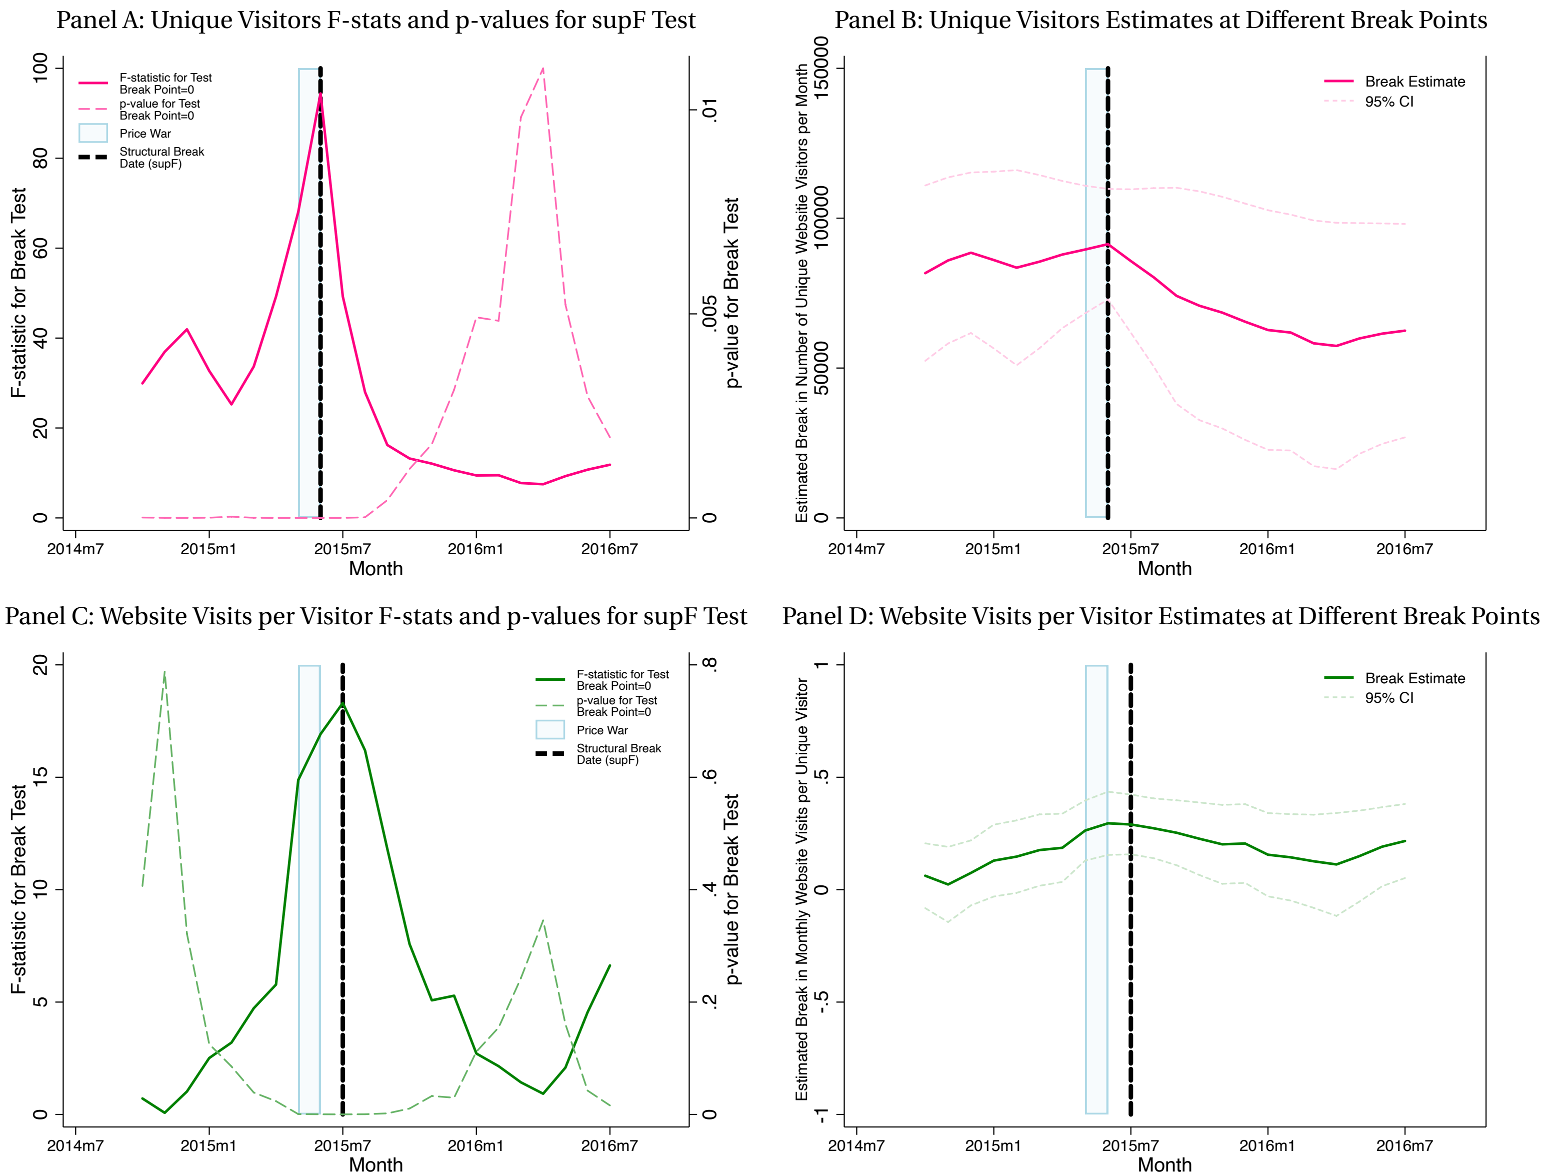

B Evolution of price dispersion after the price war

This appendix analyzes how daily price dispersion and search evolve after the May 28-June 15 price war in 2015. In particular, Panel A of Figure B.1 on the next page reproduces panel C of Figure 4 in the paper. It highlights a gradual post-war decline in daily price dispersion, which includes a downward shift. This Appendix studies what generates the shift.

B.1 When the break in price dispersion occurs

Formally, we can identify the timing of the post-war downward shift using the supF structural break test from Andrews (1993). See Appendix A above for a detailed discussion of the supF test and how we present its results. To focus on the post-war period, we conduct the test using the June 16, 2015 to July 1, 2016 post-war sub-sample. Panel (B) of Figure B.1 reports the F-statistics and p-values underlying the supF test results. We indeed find a post-war break in price dispersion in the third week of November 2015.

B.2 Decomposing price cycles into jumps and cuts

To understand what drives the November 2015 shift in price dispersion it is helpful to break down station-level and market-level price cycles into *jumps* and *cuts*. We use a set of definitions from Byrne and de Roos (2019) for classifying these parts of the cycle which we re-produce here:

Definition 1.

(i) A *station-level price jump* occurs at station i on date t if $\Delta p_{it} \geq 6$ cpl, where p_{it} is the retail price and $\Delta p_{it} = p_{it} - p_{it-1}$.

(ii) A *station-level price cycle* starts at station i on date t if $\Delta p_{it} \geq 6$ cpl. This is denoted as "day 1" of the station-level cycle. Days $2, 3, 4\ldots$ of the station-level cycle correspond to the undercutting phase, which continues until the next station-level price jump occurs and a new cycle starts. *Station-level cycle length* is the number of days between station-level price jumps.

(iii) A *market price jump* occurs on date t if $\text{median}_t(\Delta p_{it}) \geq 6$ cpl, where $\text{median}_t(\Delta p_{it})$ is the median of Δp_{it} across all stations on date t.

(iv) A *market cycle* commences on date t if $\text{median}_t(\Delta p_{it}) \geq 6$ cpl. This is denoted as "day 1" of the market cycle. Days $2, 3, 4\ldots$ of the market cycle correspond to the undercutting phase, which continues until the next market price jump occurs and a new cycle begins. *Market cycle length* is the number of days between market price jumps.

Figure B.1: Timelines of Events with Price Dispersion and Search

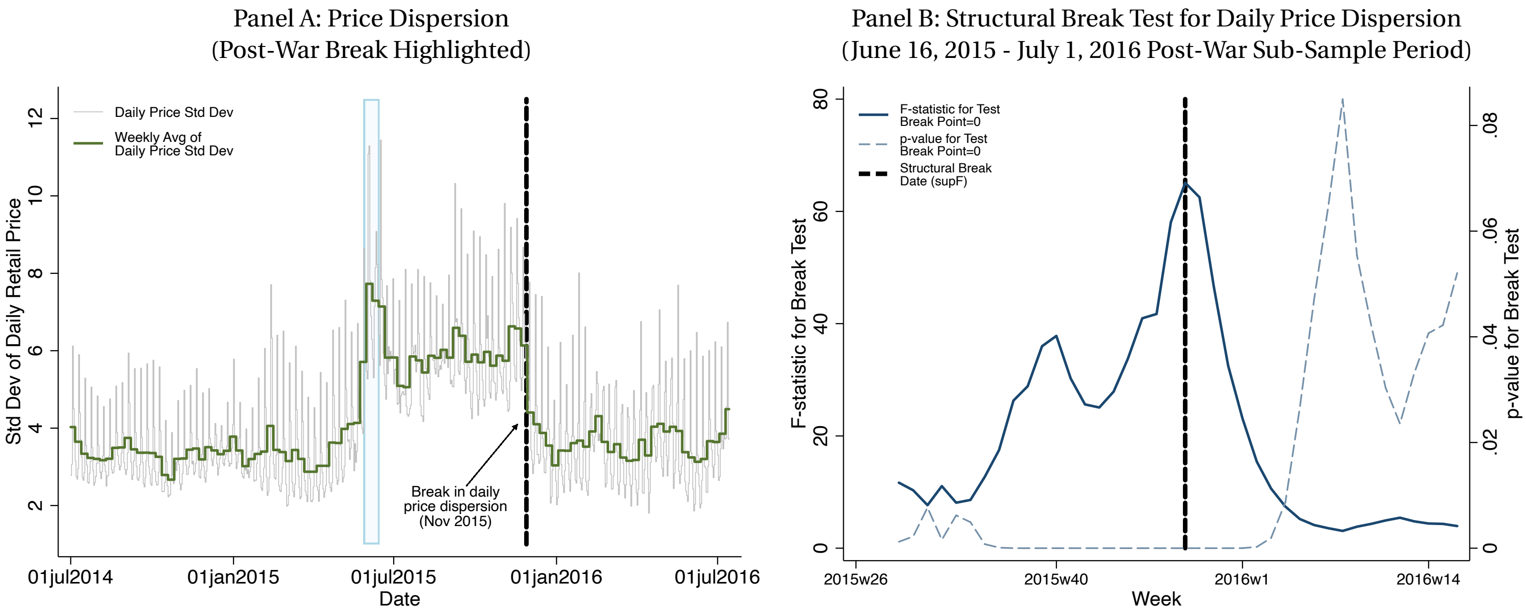

All of our results are robust to using alternatives to the 6 cpl threshold for identifying price jumps. Moreover, we focus on the four major retailers – BP, Caltex, Woolworths, and Coles – who in 2016 respectively operate 66, 49, 44, and 51 stations in the market at the start of the sample, or 210 (71%) of all 298 stations in the market at the start of the sample.[4] In Byrne and de Roos (2019) we show that BP is the historical price leader in determining the timing of price jumps and the magnitude of jumps and cuts. The May 28-June 15, 2015 price war that the paper centers on temporarily disrupts coordination on the BP-led price cycle.

B.3 Evolution of the timing and magnitude of price jumps and cuts

Figure B.2 contains three panels that together explain what drives the November 2015 shift in price dispersion. Panel A plots, by retailer, the number of stations that engage in a price jump on each Tuesday between July 1, 2015 and July 1, 2016 (i.e., after the price war). The figure shows that BP, Caltex, Woolworths, and Coles engage in station-level price jumps with nearly all of their stations on each Tuesday over this period. The one notable exception is Tuesday, January 26, 2016 when only 4 BP stations engage in a price jump. This corresponds to the national holiday Australia Day.[5]

[4]At the end of the sample the retailers operate 66, 68, 25, and 53 stations, or 212 (68%) of 311 stations in total.

[5]We do not delve into why BP engages in this behavior, but note that it is consistent with a history of BP experimentation and price signaling in the market prior to 2015 that we document in Byrne and de Roos (2019). Panel (D) of Figure B.2 shows that BP engages in a price jump with its station network on Wednesday, January 27, 2016 to reestablish coordination on the price cycle that week. This shock to coordination does, however, induce a

Caltex's steeper cycle

Panels B and C of Figure B.2 reveal the main driver of the November 2015 shift in price dispersion. Panel B plots average daily station-level price changes on Tuesdays between July 1, 2015 and July 1, 2016 by retailer. Following the price war, Caltex's stations' average daily price jump is between 20 and 25 cpl between July 1, 2015 and January 1, 2016. In contrast, average price jumps for BP, Woolworths, and Caltex stations are between 5 and 20 cpl over this period.

However, starting in November 2015, panel B further shows that average daily Tuesday price jumps for Caltex stations start converging toward those of its rivals. From this point until March 15, 2016, average station-level price jumps on Tuesdays are similar in magnitude across the four major retailers. Then, between March 15 and May 15 in 2016, Caltex, Woolworths, and Coles begin coordinating on higher Tuesday price jumps between 22 and 27 cpl, while BP's average jump remains between 12 and 17 cpl. After May 15, 2016, however, we see the four retailers again coordinate on the same average price jump between 10 and 15 cpl.

Importantly, this evolution in coordination on price jump magnitudes occurs in tandem with a related evolution in daily price changes on Wednesdays. We illustrate this by plotting daily, non-Tuesday average station-level price changes in Panel C of Figure B.2. Before November 2015, there is much less coordination on price changes on other days of the week between Caltex and its rivals. In particular, the large negative blue spikes in panel C correspond to Caltex engaging in roughly 10 cpl price cuts on Wednesdays, which compares to 2 to 5 cpl price cuts by its rivals. After November 2015, Caltex stops engaging in large price cuts on Wednesdays as it coordinates on Tuesday price jumps with its rivals; this is highlighted in panel C of Figure B.2. This coordination on non-Tuesday price cuts persists throughout the remainder of the sample, with the exception of a four-week period in May 2016 where Caltex and Coles coordinate on large Wednesday price cuts after engaging in large Tuesday price jumps over this period.

In sum, our findings in panels B and C of Figure B.2 lead us to our main driver of the November 2015 shift in daily market-level price dispersion: *Caltex engages in steeper price cycles compared to its rivals* between July 1, 2015 and November 2015, and then stops.

temporary search response on the FuelWatch platform, which can be seen in Panel D of Figure 4.

Figure B.2: Timing and Magnitude of Daily Station-Level Price Jumps and Cuts by Retailer

Panel A: No. of Station-Level Price Jumps by Retailer on Tuesdays

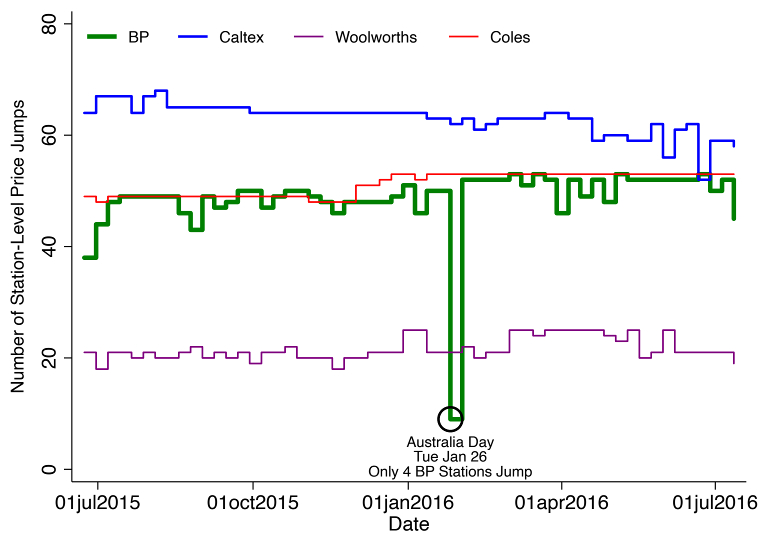

Panel B: Avg. Daily Station-Level Price Change by Retailer on Tuesdays

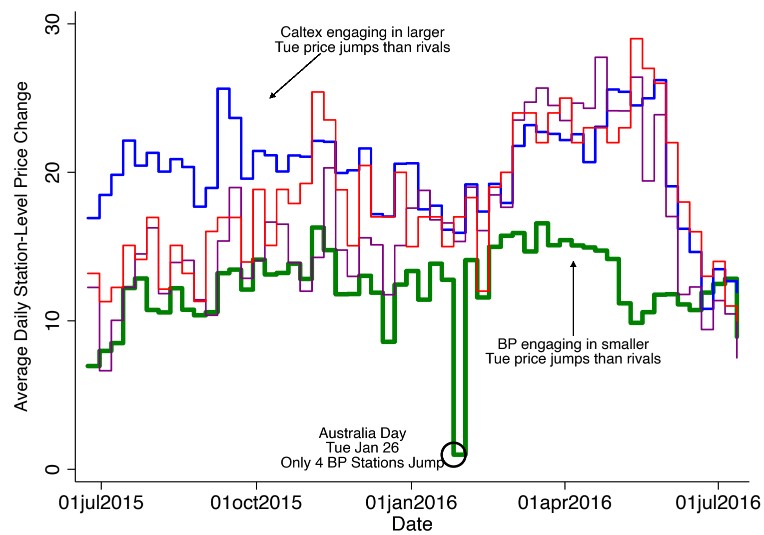

Panel C: Avg. Daily Station-Level Price Change by Retailer Excl. Tuesdays

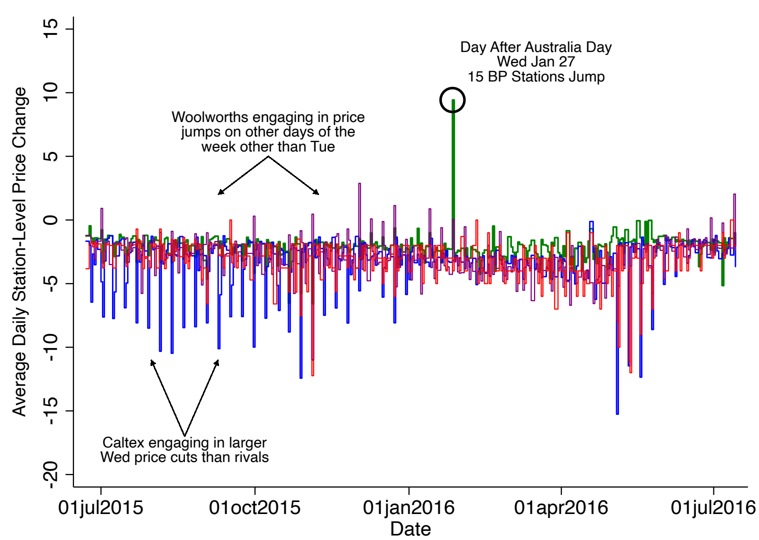

C Details on constructing consideration sets

We construct the distribution of consideration sets across consumers in the city, P_J, in four steps. First, we obtain from the Australian Bureau of Statistics (ABS) (https://www.abs.gov.au/) an origin-destination (O-D) table which was constructed as part of the 2016 Census. Origin locations are narrowly-defined ABS Statistical Area 1 (SA1) census blocks which contain 400 people on average. Destination locations correspond to ABS Place of Work Destination Zones (DZN) which in our sample contain 1350 people on average. In total, there are 4,250 SA1s and 593 DZNs in the market. The O-D table contains the joint distribution of households that commutes from each SA1 to each DZN for work day-to-day. We obtained this table along with the corresponding GIS mapping files for the SA1s and DZNs from the ABS for our analysis.

In the second step, we obtain the universe of search results that the FuelWatch platform can potentially return. Search on the platform occurs at the postcode–level, which is a higher level of aggregation than the SA1 or DZN level. With the list of all postcodes in the market obtained from the ABS, we scrape the search results from the FuelWatch platform by inputting each postcode into the search engine and recording the list of individual stations that a given search returns. After scraping the website, we are left with a list of postcodes, station names, and addresses that are returned for each postcode–based search.

Third, using ArcGIS we conduct a spatial join that links each SA1 and DZN to postcodes that they overlap with. For each SA1 and DZN this allows us to construct the list of stations that a FuelWatch search returns based on the postcode a given SA1 or DZN lies within. If an SA1 or DZN intersects with more than one postcode, then this list contains the set of unique stations obtained by searching all the postcodes that an SA1 or DZN overlaps with.[6]

Finally, we again use ArcGIS to construct optimal (fastest) commuting paths from each SA1 to each DZN based upon the market's entire road network and each road segment's speed limit. For each SA1–DZN pair we find the list of stations that sit along the path. Having detailed station address data from the FuelWatch website is key to identifying stations along each SA1–DZN commuting route in the market.

After conducting these four steps, we obtain three lists of potentially searched stations for a given O-D pair: those who are listed from a FuelWatch search at origin (O), destination (D), and along the O-D route (R). We assume that the consideration set of consumer type τ_i, J_i, are all the stations in O, D, and R. The distribution of consideration sets across households, P_J, is the distribution of the driving population across O-D pairs from the 2016 Census O-D table.

[6] In the case where a given SA1 or DZN overlaps with multiple postcodes, none of our results are changed if we instead match SA1s or DZNs uniquely to a postcode based upon the postcode that the majority of a given SA1's or DZN's spatial area lies within.

D Robustness checks on model specification

In this appendix, we investigate the implications of alternative specifications of search benefits and costs for our analysis in Section 4 of the paper. Recall that in our model, we specify the gains to search for consumer i at time t as y_{it}, where

$$y_{it} = 1\{\max\{g_{it}, g_{it+1}\} > s_{it}\},$$
$$g_{it} = (\overline{p}_{it} - \min\{\mathbf{p}_{it}\}) \times \kappa_i.$$

Further recall search costs s_{it} are specified as

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

where c_i and f_i are the consumer's recurrent and startup search costs. Thus, we assume that consumers are concerned about the current day's benefits of search, incorporating the potential to purchase today or tomorrow, and their evaluation of those search benefits is based on current information about the price distribution. Moreover, in computing g_{it} we assume i's consideration set of stations consists of those around their home (origin), at their location of work (destination), and along their commuting route from home to work.

We consider four alternative modeling choices in this appendix. Section D.1 considers variation in the forward-looking horizon over which consumers evaluate search benefits. In Section D.2, we evaluate the implications of backward-looking, adaptive expectations for computing current expected search benefits. Section D.3 checks robustness of our recurrent and startup search cost estimates to different combinations of origin/destination/route combinations in constructing consideration sets. Finally, in Section D.4 we adopt the Log Normal distribution for recurrent and startup search costs, allow correlation between recurrent and startup search costs, introduce heterogeneity in the size of the consideration set around home and workplace, and allow for correlation between a household's recurrent and startup search costs with the size of their consideration set around their home and place of work.

D.1 Consumers with a dynamic perspective

Our model in Section 4 presumes 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

$$
\begin{aligned}
V(w_{it}) &= \max_{\chi_{it}\in\{0,1\}} \chi_{it}\left(\overline{u}-\kappa_i \min\{\mathbf{p}_{it}\} - s_{it}\right) + (1-\chi_{it})\left(\overline{u}-\kappa_i \overline{p}_{it}\right) + \delta \mathbb{E}_t V(w_{it+1}), \\
w_{it+1} &= w_{it} + (1-w_{it})\chi_{it},
\end{aligned}
$$

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.[7] The parameter δ describes the rate at which consumers discount the next fuel purchase. This could reflect impatience and concerns about the decay or obsolecence 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. 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 $u(.)$ is defined in equation (1) in the paper.

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

$$
\begin{aligned}
V(0) = \max_{\chi_{it}\in\{0,1\}} & \chi_{it}\left(\overline{u}-\kappa_i \min\{\mathbf{p}_{it}\} - c_i - f_i\right) + (1-\chi_{it})\left(\overline{u}-\kappa_i \overline{p}_{it}\right) \\
& + \delta\left(\chi_{it} V(1) + (1-\chi_{it}) V(0)\right).
\end{aligned}
$$

[7] For simplicity, we ignore the intertemporal searh opportunities presented by the FuelWatch program in this formulation, and consider search gains in period t based solely on the period t price distribution.

Consumer i searches in period t iff

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

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

$$(1-\delta)\left(\overline{u}-\kappa_i \min\{\mathbf{p}_{it}\}-c_i-f_i\right)+\delta u(c_i)>\overline{u}-\kappa_i\overline{p}_{it}. \tag{2}$$

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 $u(c_i)=\overline{u}-\kappa_i\min\{\mathbf{p}_{it}\}-c_i$. Substituting into (2) leads to the conclusion that i searches iff $g_{it}>c_i+(1-\delta)f_i$. Next, suppose that $g_{it}\le c_i$. In this case, $u(c_i)=\overline{u}-\kappa_i\overline{p}_t$. Suppose further that $\chi_{it}=1$. Substituting into (2) leads to the condition $g_{it}>c_i+f_i$, a contradiction. Therefore $\chi_{it}=0$ whenever $g_{it}\le c_i$. 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. □

Proposition 1 highlights 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 c_i. However, particularly for patient consumers, inferred startup search costs f_i 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 engage in search for the first time. 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.

D.2 Adaptive expectations

In this robustness check, we allow consumers to form adaptive expectations over search benefits, whereby their expectations depend on realized gains from search in recent weeks. Here, we abstract from weekly cyclical variation in search benefits and focus on the long-run value of adopting the search platform. To do this, we adapt our model to a weekly frequency. We redefine search intensity, $\hat{q}_{it}$, to be the weekly average of the daily share of searchers in the market. We calculate search benefits as the simple average of daily search benefits over the previous x weeks, including the current week, for $x \in \{1, 2, 3, 4\}$. Our base case, $x = 1$, is the weekly analogue of the daily model we present in the paper. In this way, we vary the horizon over which adaptive expectations are formed. We leave all other elements of the model unchanged.

Table D.1 presents estimates of search cost distribution parameters based on each horizon. Columns 1-4 contain estimates based on a 1, 2, 3, and 4 week backward-looking time horizon. The top panel contains estimates for the full model for each specification, and the bottom panel contains estimates based on a restricted model without startup search costs. Figure D.1, analogous to Panel C of Figure 6 in the paper, depicts the cumulative distribution of search costs for the restricted model, and for the full model, evaluated both before and after the price war. Panels A, B, C, and D, respectively, depict the search cost distribution for the one, two, three, and four week horizon. Figure D.2 illustrates model predictions for both the full and restricted model for each specification. Panels on the left of the figure depict the full model, and panels on the right show the restricted model. The panels are organised vertically, with the shortest horizon at the top, and the longest horizon at the bottom. Finally, Figure D.3, organized in the same manner as Figure D.1, shows the relationship between search gains and the evolution of the stock of active consumers.

Comparing Panel A of Figure D.1 with Panel C of Figure 6 in the paper suggests that our inferences about the search cost distribution are qualitatively similar with the daily and weekly models. Estimated search costs are lower using weekly data. For example, the 20$^{\text{th}}$ percentile startup cost is \$12.67/day in the weekly model, compared to \$15.70/day in the daily model. This is because the weekly model smoothes over weekly variation in search intensity and benefits, understating peaks in daily search activity. Comparing panels A through D of Figure D.1 suggests inferences about the search cost distribution are unaffected by the choice of backward-looking horizon.

As we found with the daily model, Table D.1 and Figure D.2 reveal substantial differences in the predictive ability of the full and restricted models. The objective function, $\hat{G}(\theta)$, is an order of magnitude greater in the model without startup costs. Moreover, for each horizon, the full model captures the permanent shift in search intensity in the data, while the restricted model predicts only a temporary shift in search intensity. The impact of the horizon on model predic-

Table D.1: Estimation Results, Weekly Model, Backward Looking Expectations

Model	Current week	Two weeks	Three weeks	Four weeks
		WITH STARTUP COSTS		
Recurrent costs				
μ_c	0.599	0.200	0.207	0.399
	(0.217)	(0.029)	(0.036)	(0.065)
σ_c	2.100	30.639	20.609	14.443
	(0.583)	(2.674)	(4.349)	(2.001)
Startup costs				
μ_f	5.036	2.991	2.995	1.989
	(0.170)	(0.193)	(0.130)	(0.146)
σ_f	4.061	7.376	7.419	12.112
	(0.242)	(0.590)	(0.348)	(1.057)
Objective, $G(\hat{\theta})$	0.011	0.011	0.011	0.012
		NO STARTUP COSTS		
Recurrent costs				
μ_c	0.787	0.613	0.864	0.788
	(0.022)	(0.013)	(0.032)	(0.026)
σ_c	86.186	195.16	65.527	85.607
	(7.091)	(10.262)	(6.253)	(7.488)
Objective, $G(\hat{\theta})$	0.112	0.098	0.106	0.098

Notes: Robust standard errors are in parentheses (). The number of observations is $T = 104$ dates. All calculations assume consumers purchase 50 liters of gasoline. Model indicates backward-looking forecast horizon for households' adaptive expectations.

tions can most easily be seen by examining the period surrounding the price war in the middle of 2015. In panel A of Figure D.2, with search benefits evaluated using prices from the current week only, the transition in predicted search intensity is rapid, with a sharp jump in predicted search occuring in a single week. By contrast, in panel G, search benefits are evaluated over the previous four weeks. In this case, search intensity is predicted to rise more gradually in response to the price war as consumers gradually revise upwards their expectations of search benefits.

We observe similar patterns in Figure D.3. In all treatments, the stock of active consumers increases with the unprecedented benefits of search during the price war in the middle of the sample. However, as expected, the pace of uptake is more gradual the longer is the backward-looking horizon of expectation formation.

Figure D.1: Search cost distribution

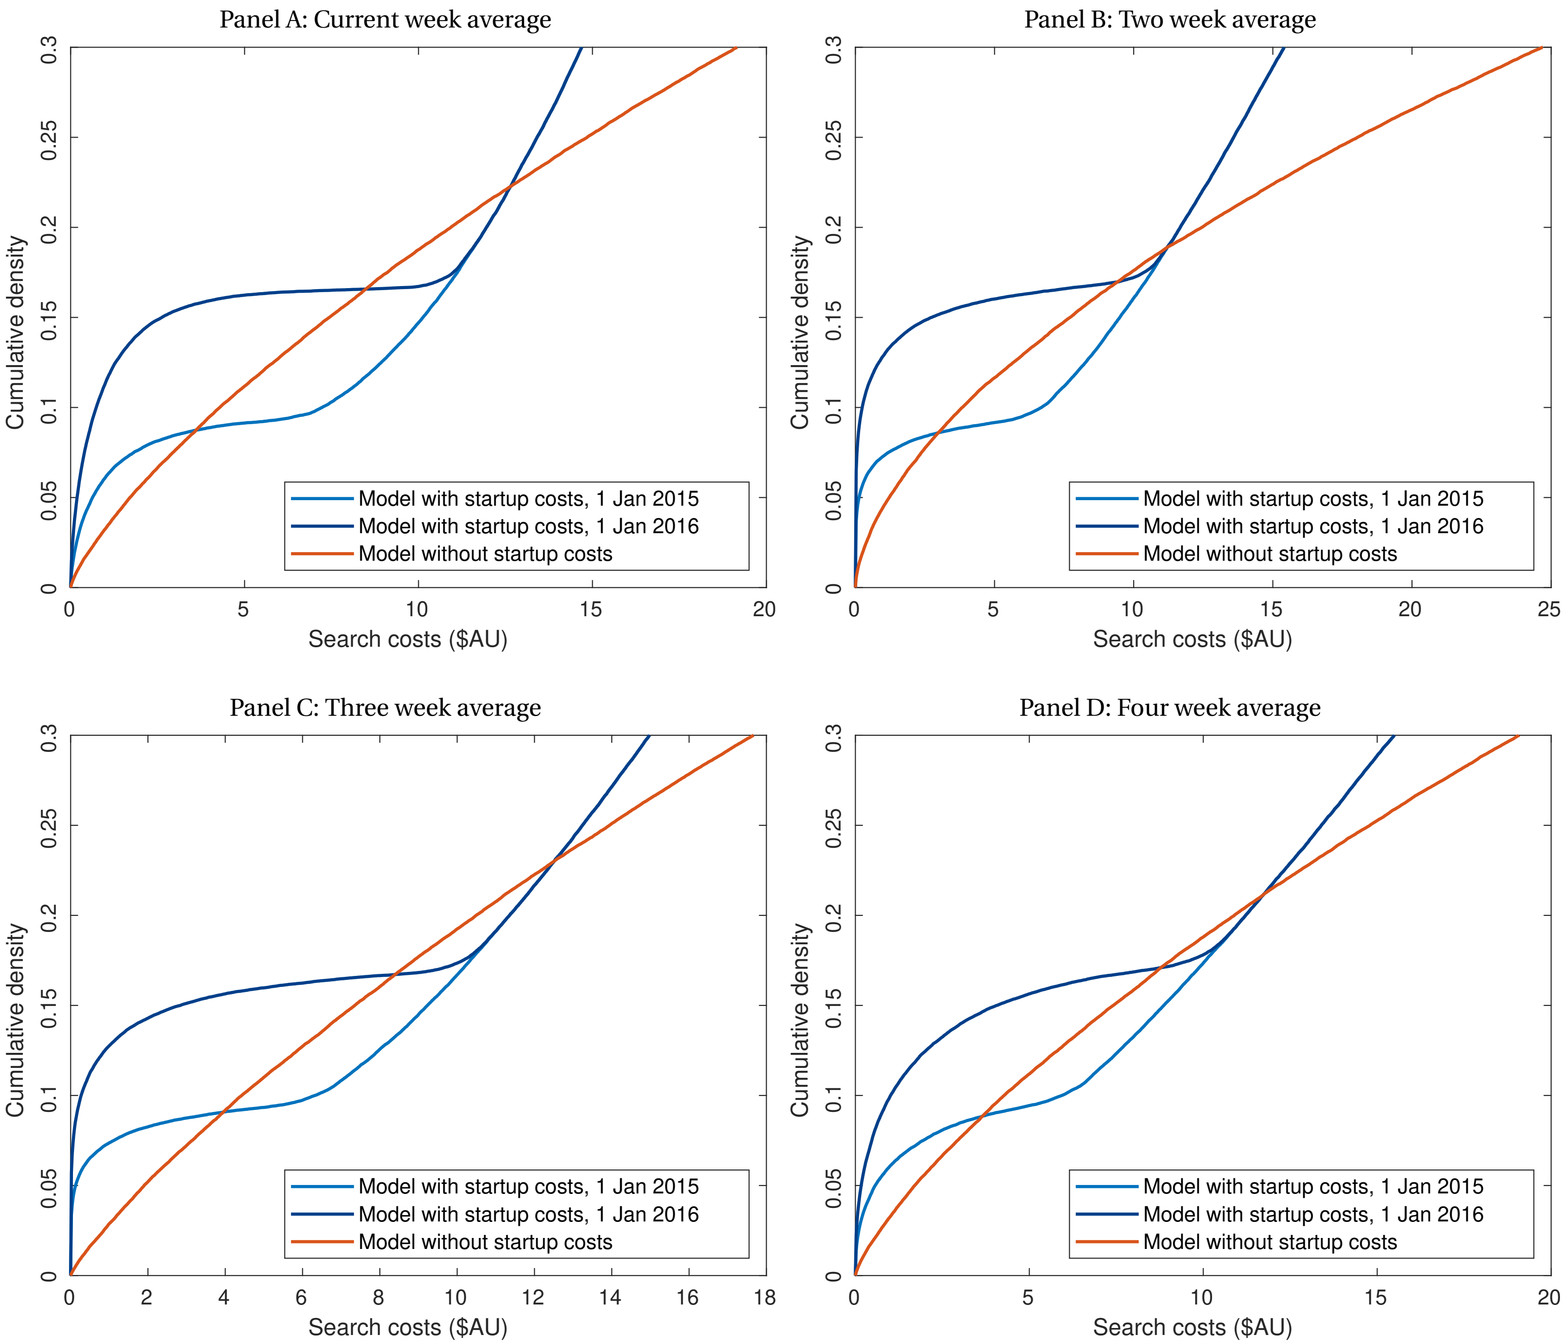

Figure D.2: Model predictions

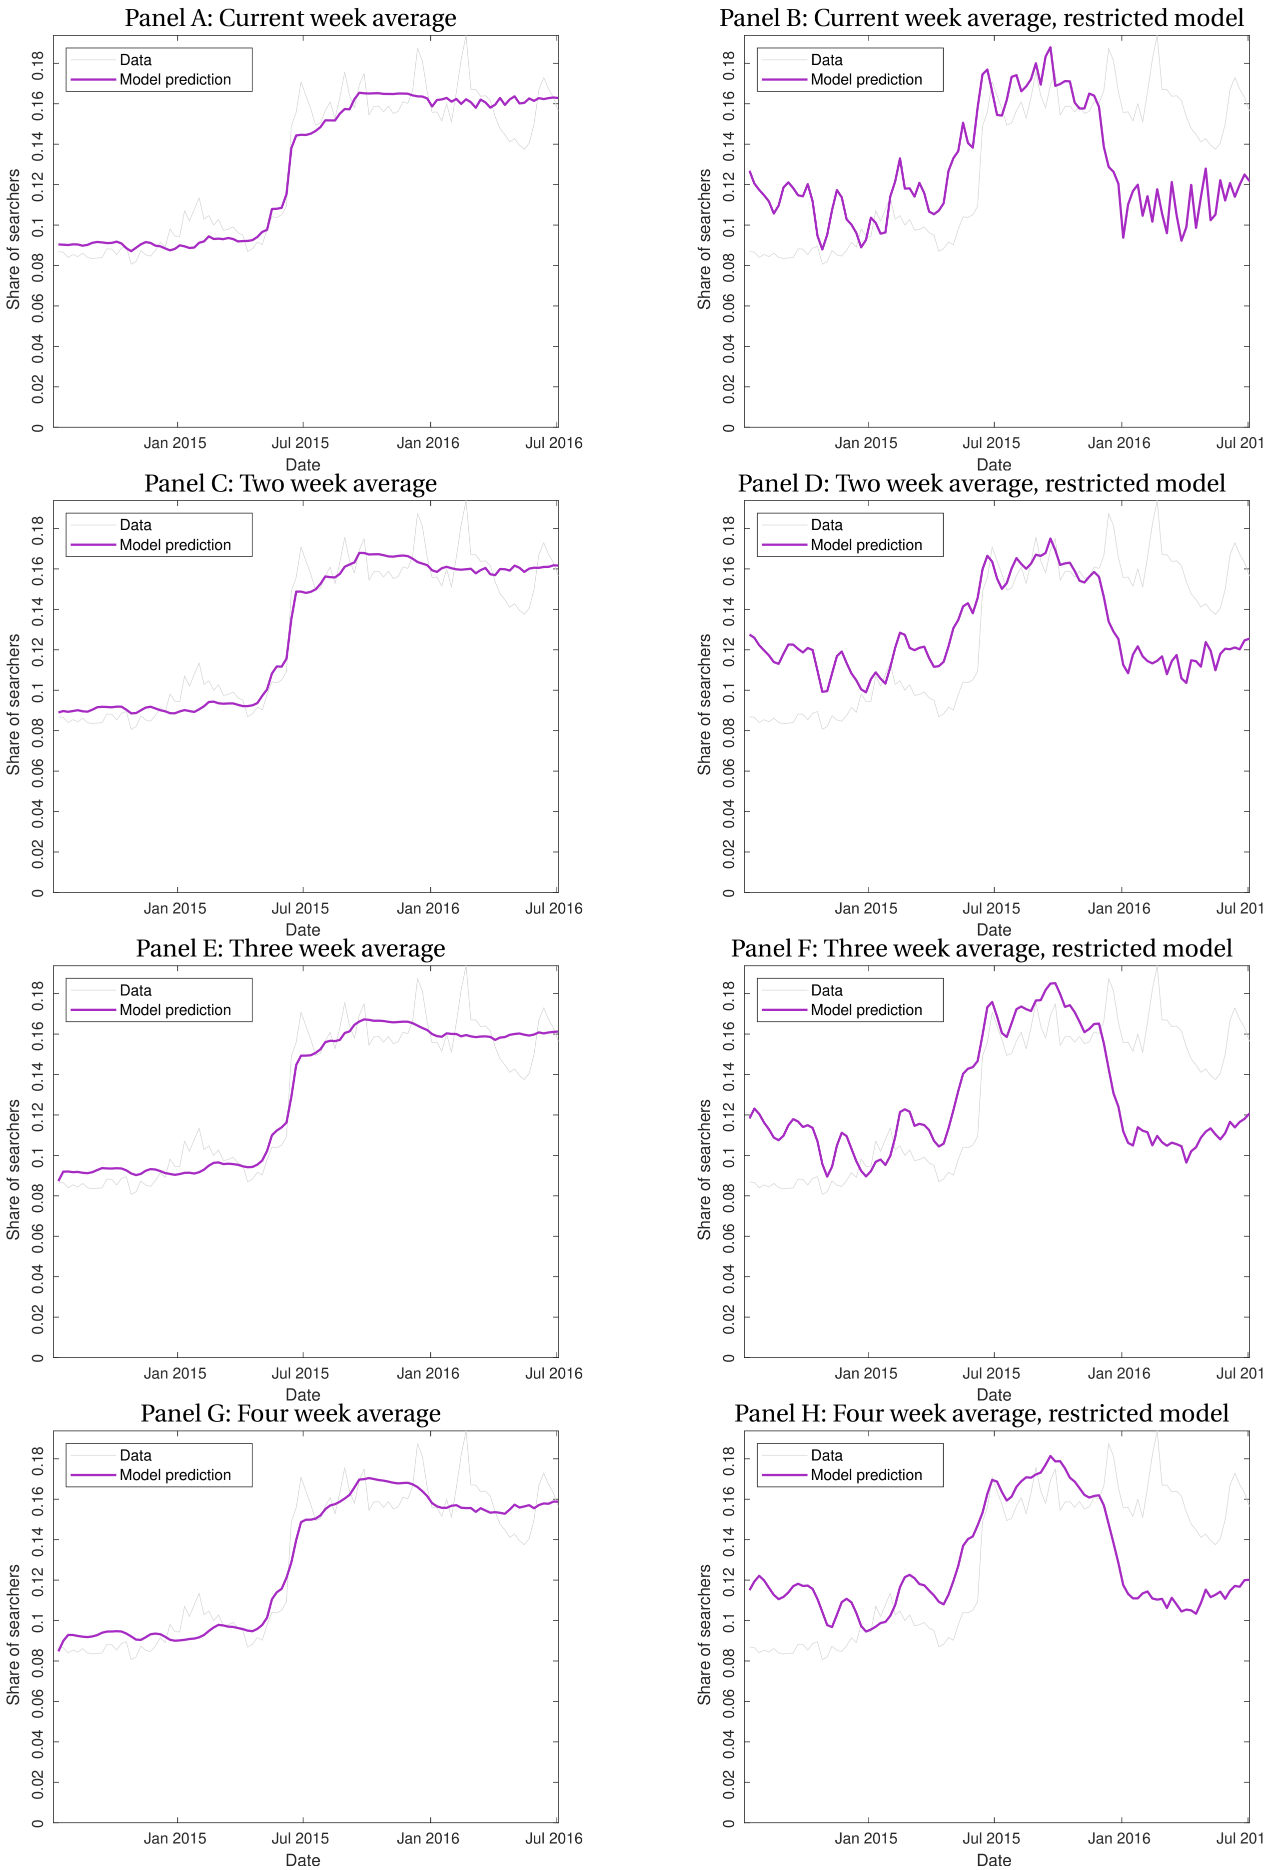

Figure D.3: Gains From Search

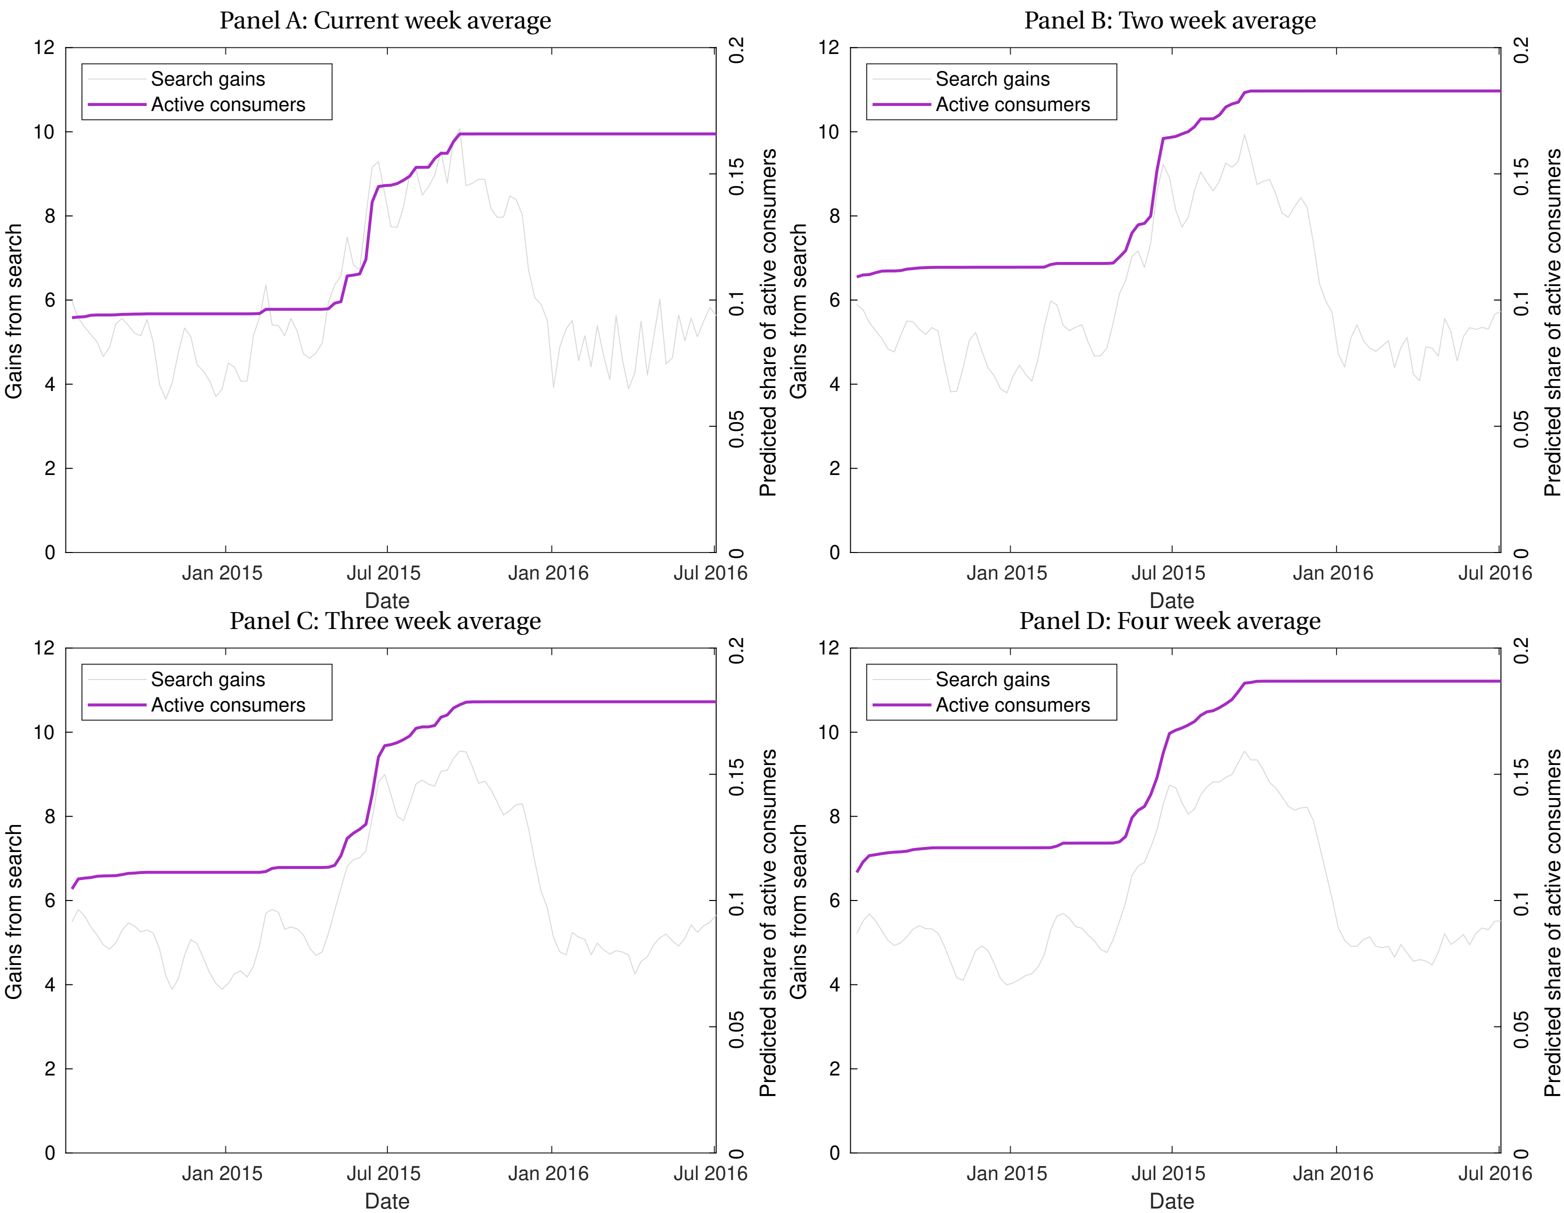

D.3 Alternative consideration sets

In the paper, we assumed that the consideration set of consumer i was drawn from stations near here home (origin, O), her place of work (destination, D) and along here O-D travel route (R). Here, we consider alternative constructions of the consideration set. Variation in the size of the consideration set leads to variation in the benefits of search, and therefore variation in the estimated magnitude of search costs.

Table D.2 contains our consideration set size robustness checks. Columns 1–6 of the table contain, respectively, specifications with consideration sets drawn from O only, D only, the origin only (O), the destination only (D), the route only (R), the origin and route (O+R), the destination and route (D+R), and origin and destination (O+D).

Columns 7 and 8 contain hybrid specifications in which the consideration set of the consumer depends on whether she uses the search platform or not. In the event that she uses the search platform, her consideration set is as specified in Section 4. However, if she chooses not to use the platform, we allow her consideration set to be narrower. In column 7 (H1), she considers only stations on her travel route if she does not use the search platform. In column 8 (H2), she considers stations on her travel route, and only the stations in the same suburb as her origin and destination, but not neighboring suburbs as in Section 4 from the paper. For each specification, we present estimates based on the full model in the top panel, and estimates based on the restricted model in which we abstract from startup search costs in the lower panel.

Finally, Figure D.4, analogous to Panel C of Figure 6, presents the estimated cumulative distribution of search costs for the restricted model, and the full model, both before and after the price war. Panels A to H correspond to the specifications in columns 1 to 8 of Table D.2. Figures D.5 and D.6, containing predictions for platform search over the sample period for the full and restricted model for each specification, are similarly organized. Figure D.7, illustrating the time profile of platform adoption in the full model, is also organized in the same manner.

With the partial exception of the model with a consideration set drawn from the travel route alone (Model R of Table D.2, Panel C of Figures D.4, D.5, D.6 and D.7), our estimates share the same qualitative features as our primary specification. Consider first the distribution of search costs shown in Figure D.4. For each specification, the model with startup costs contains an inflection in the distribution of search costs as the mix of consumers adjusts: consumers to the left of the distribution have incurred their startup costs, and consumers to the right of the distribution have not. Further, the change in the distribution of search costs can be seen in the upward shift in the distribution to the left of the figure. From Table D.2, we can see that model fit is substantially improved in each case, once startup costs are incorporated. Figures D.5 and D.6 illustrate that, for each specification, the model can only explain the permanent increase in search activity following the price war once startup costs are included.

Table D.2: Estimation Results with Alternative Consideration Sets

	Consideration Set							
	O	D	R	O+R	D+R	O+D	H1	H2
	WITH STARTUP COSTS							
Recurrent costs								
μ_c	0.234	0.374	0.212	0.210	0.219	0.255	0.213	0.213
	(0.015)	(0.022)	(0.070)	(0.017)	(0.015)	(0.017)	(0.016)	(0.015)
σ_c	39.685	22.946	2.103	36.336	40.568	40.085	32.963	40.393
	(3.311)	(2.167)	(1.148)	(3.547)	(2.873)	(3.415)	(3.271)	(3.330)
Startup costs								
μ_f	3.021	3.020	1.938	5.982	4.022	3.925	7.134	5.162
	(0.080)	(0.125)	(0.059)	(0.194)	(0.171)	(0.147)	(0.142)	(0.210)
σ_f	8.694	8.747	10.823	4.009	6.049	6.447	3.708	5.064
	(0.413)	(0.612)	(0.716)	(0.118)	(0.360)	(0.322)	(0.144)	(0.246)
Objective, $G(\hat{\theta})$	0.531	0.552	0.553	0.554	0.533	0.543	0.536	0.543
	WITHOUT STARTUP COSTS							
Recurrent costs								
μ_c	0.534	0.461	0.463	0.418	0.534	0.505	0.501	0.484
	(0.005)	(0.004)	(0.010)	(0.003)	(0.005)	(0.005)	(0.004)	(0.004)
σ_c	239.11	582.82	57.063	992.33	255.99	390.82	447.75	505.13
	(7.210)	(12.652)	(4.203)	(16.450)	(7.750)	(11.442)	(11.421)	(11.912)
Objective, $G(\hat{\theta})$	1.241	1.269	1.182	1.237	1.271	1.299	1.347	1.287

Notes: Robust standard errors are in parentheses (). The number of observations is $T = 731$ dates. All calculations assume consumers purchase 50 liters of gasoline. Model indicates consideration set for search: O: Origin; D: Destination; R: on Route; O+R: Origin and Route; D+R: Destination and Route; O+D; Origin and Destination; H1: Searchers consider Origin, Desination, and Route, Non-searchers consider Route; H2:Searchers consider Origin, Desination, and Route, No-searchers consider Route and local suburb only for Origin and Destination.

The main differences between specifications follow from differences in the size of the consideration sets. In our baseline specification, stations from the consumer's origin, destination, and travel route are included in the consideration set. The average consumer has 34.72 stations in her consideration set (s.d. 14.77). By contrast, when the consideration set is restricted to stations on a consumer's travel route alone (Model R), the average consideration set contains 2.31 stations (s.d. 1.35). This leads to lower search benefits in Model R, and therefore also lower estimated search costs: in the baseline model, 20$^{\text{th}}$ percentile startup costs are \$15.70 and median recurrent costs are \$0.93; in Model R, 20$^{\text{th}}$ percentile startup costs are \$8.20 and median recurrent costs are \$0.02.

Comparing Panel C of Figures D.5 and D.6 to the other models reveals some subtle differences in model predictions. When the consideration set is defined more restrictively (as in Panel

C), 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 (Figure D.5) and the restricted model without startup search costs (Figure D.6). Finally, Figure D.7 also suggests that, when the consideration set is more restrictive, to rationalize the volume of search activity, a greater proportion of consumers are predicted to incur their startup search costs.

Figure D.4: Search cost distribution

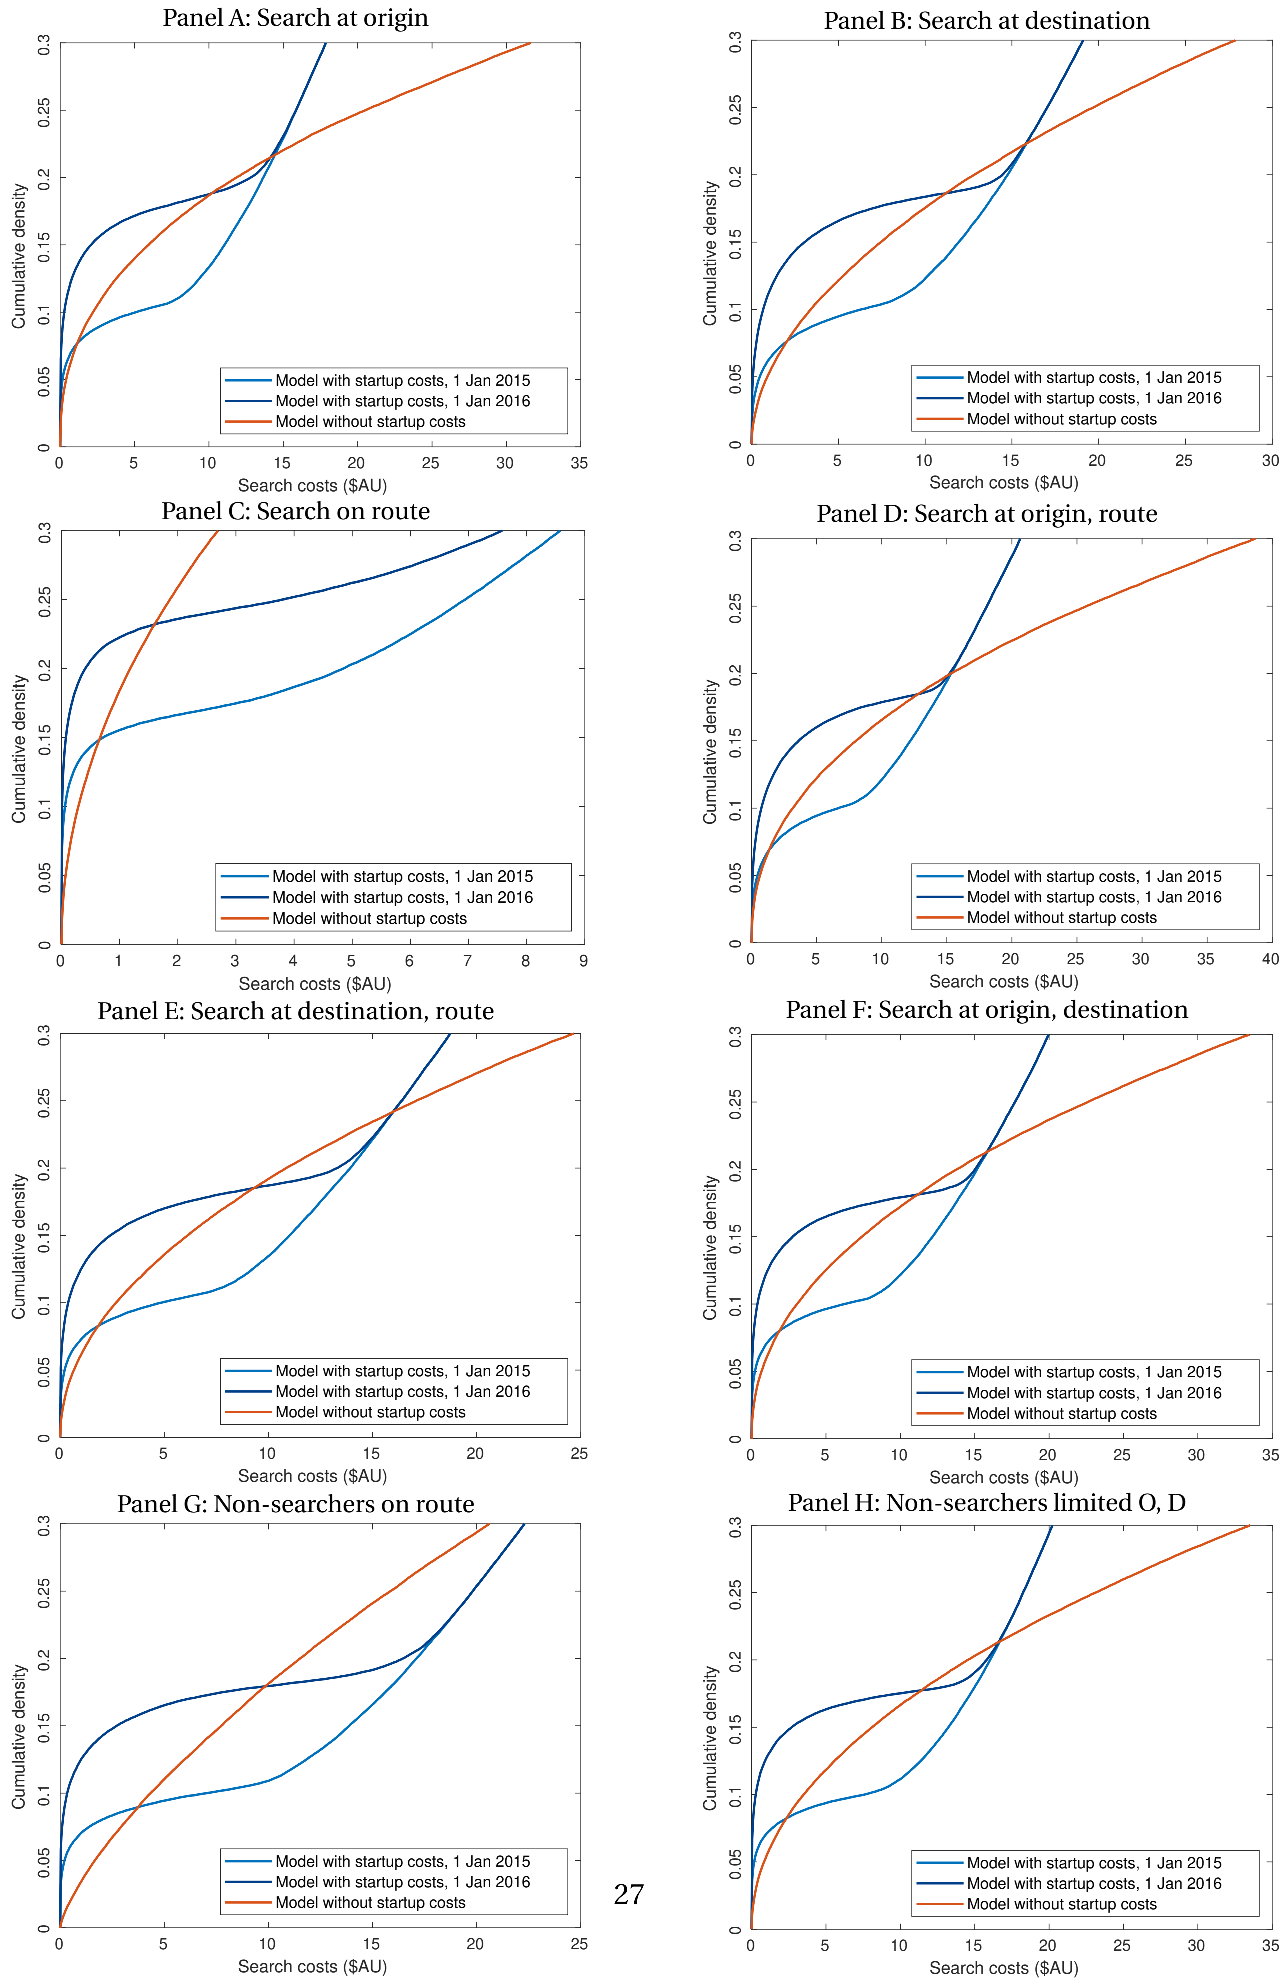

Figure D.5: Model predictions

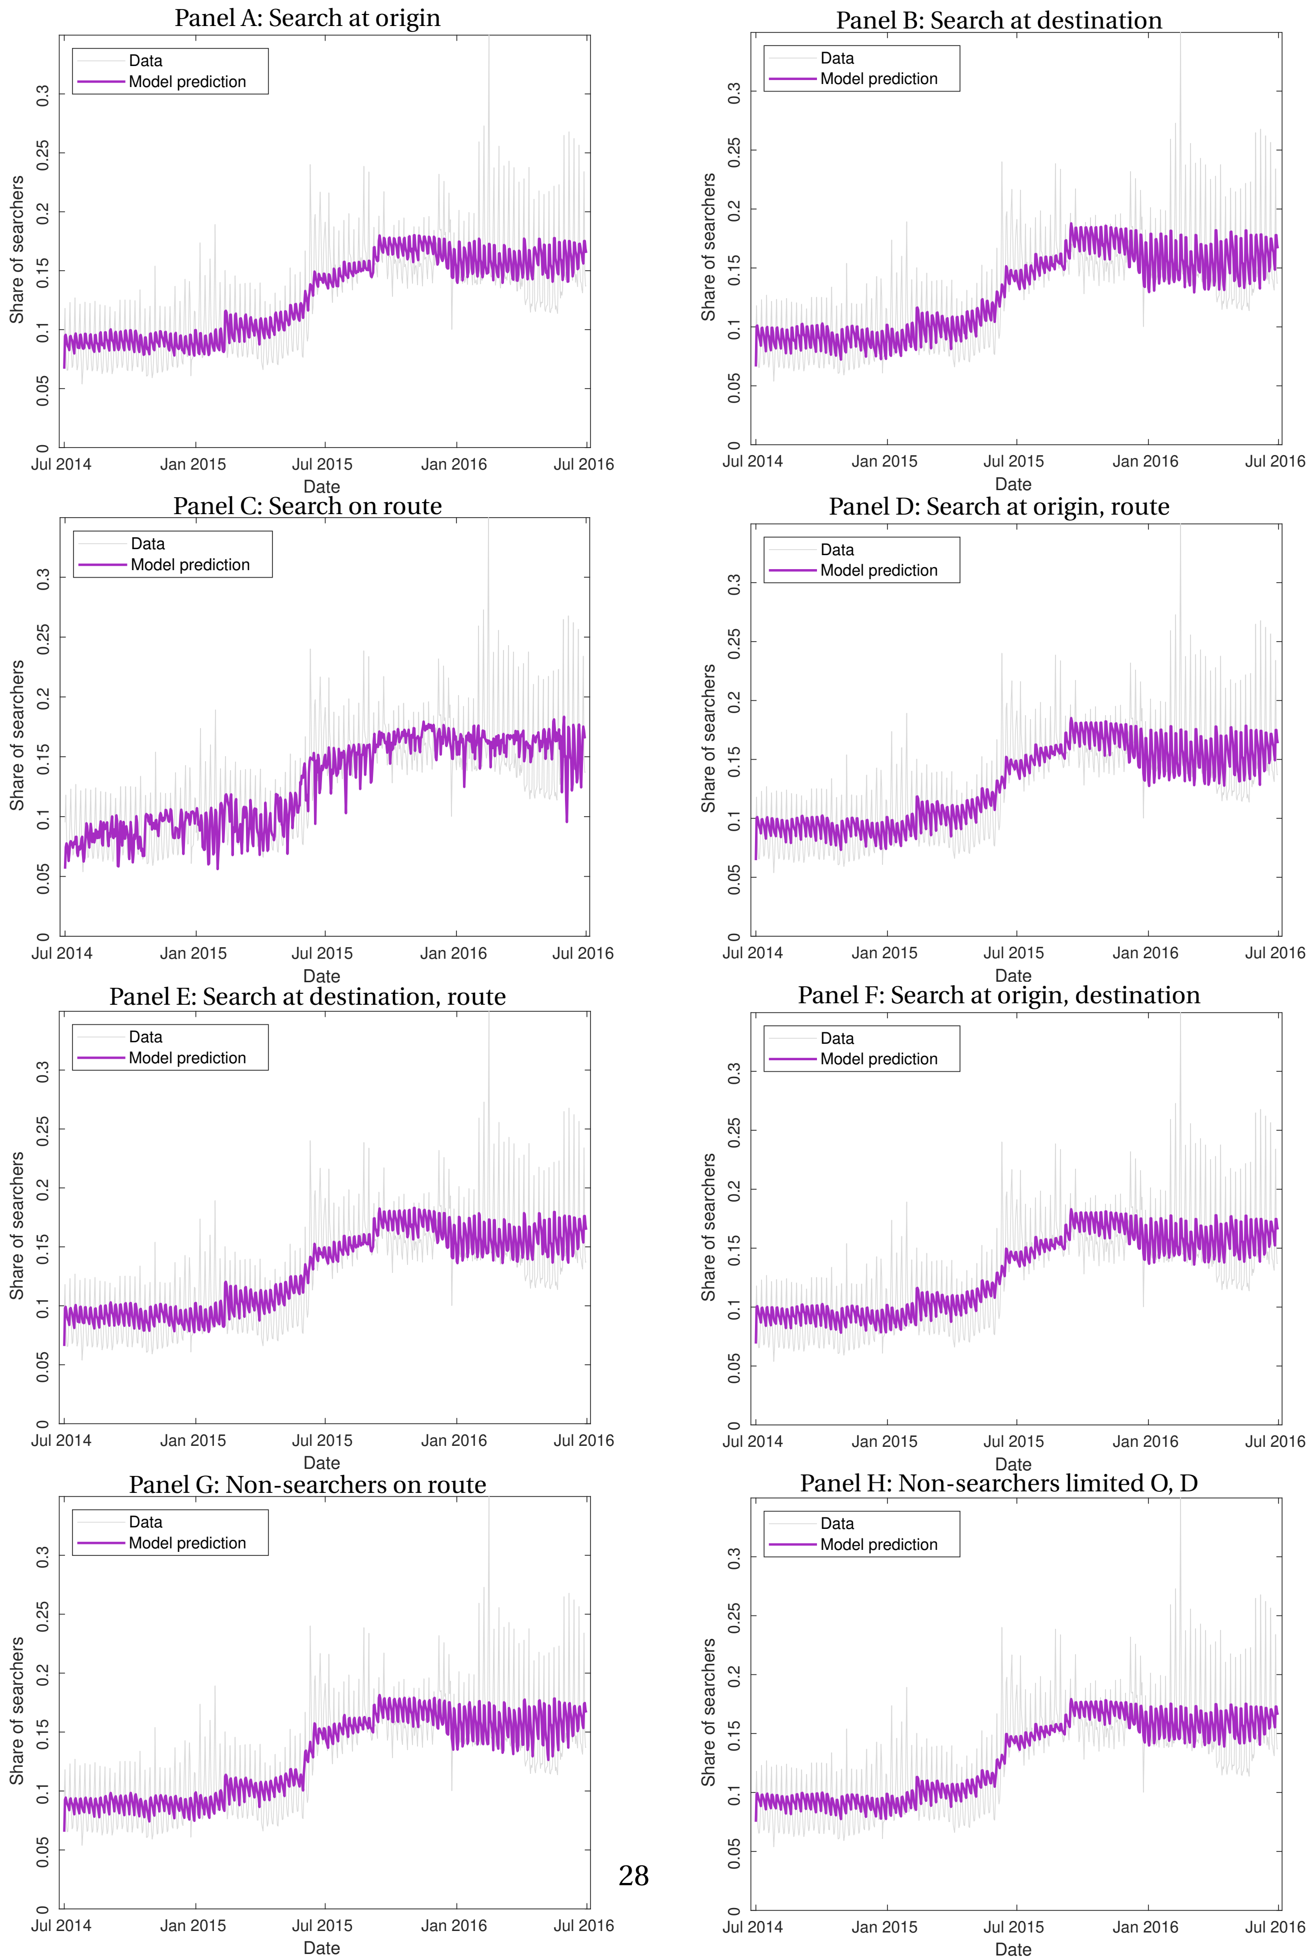

Figure D.6: Restricted model predictions

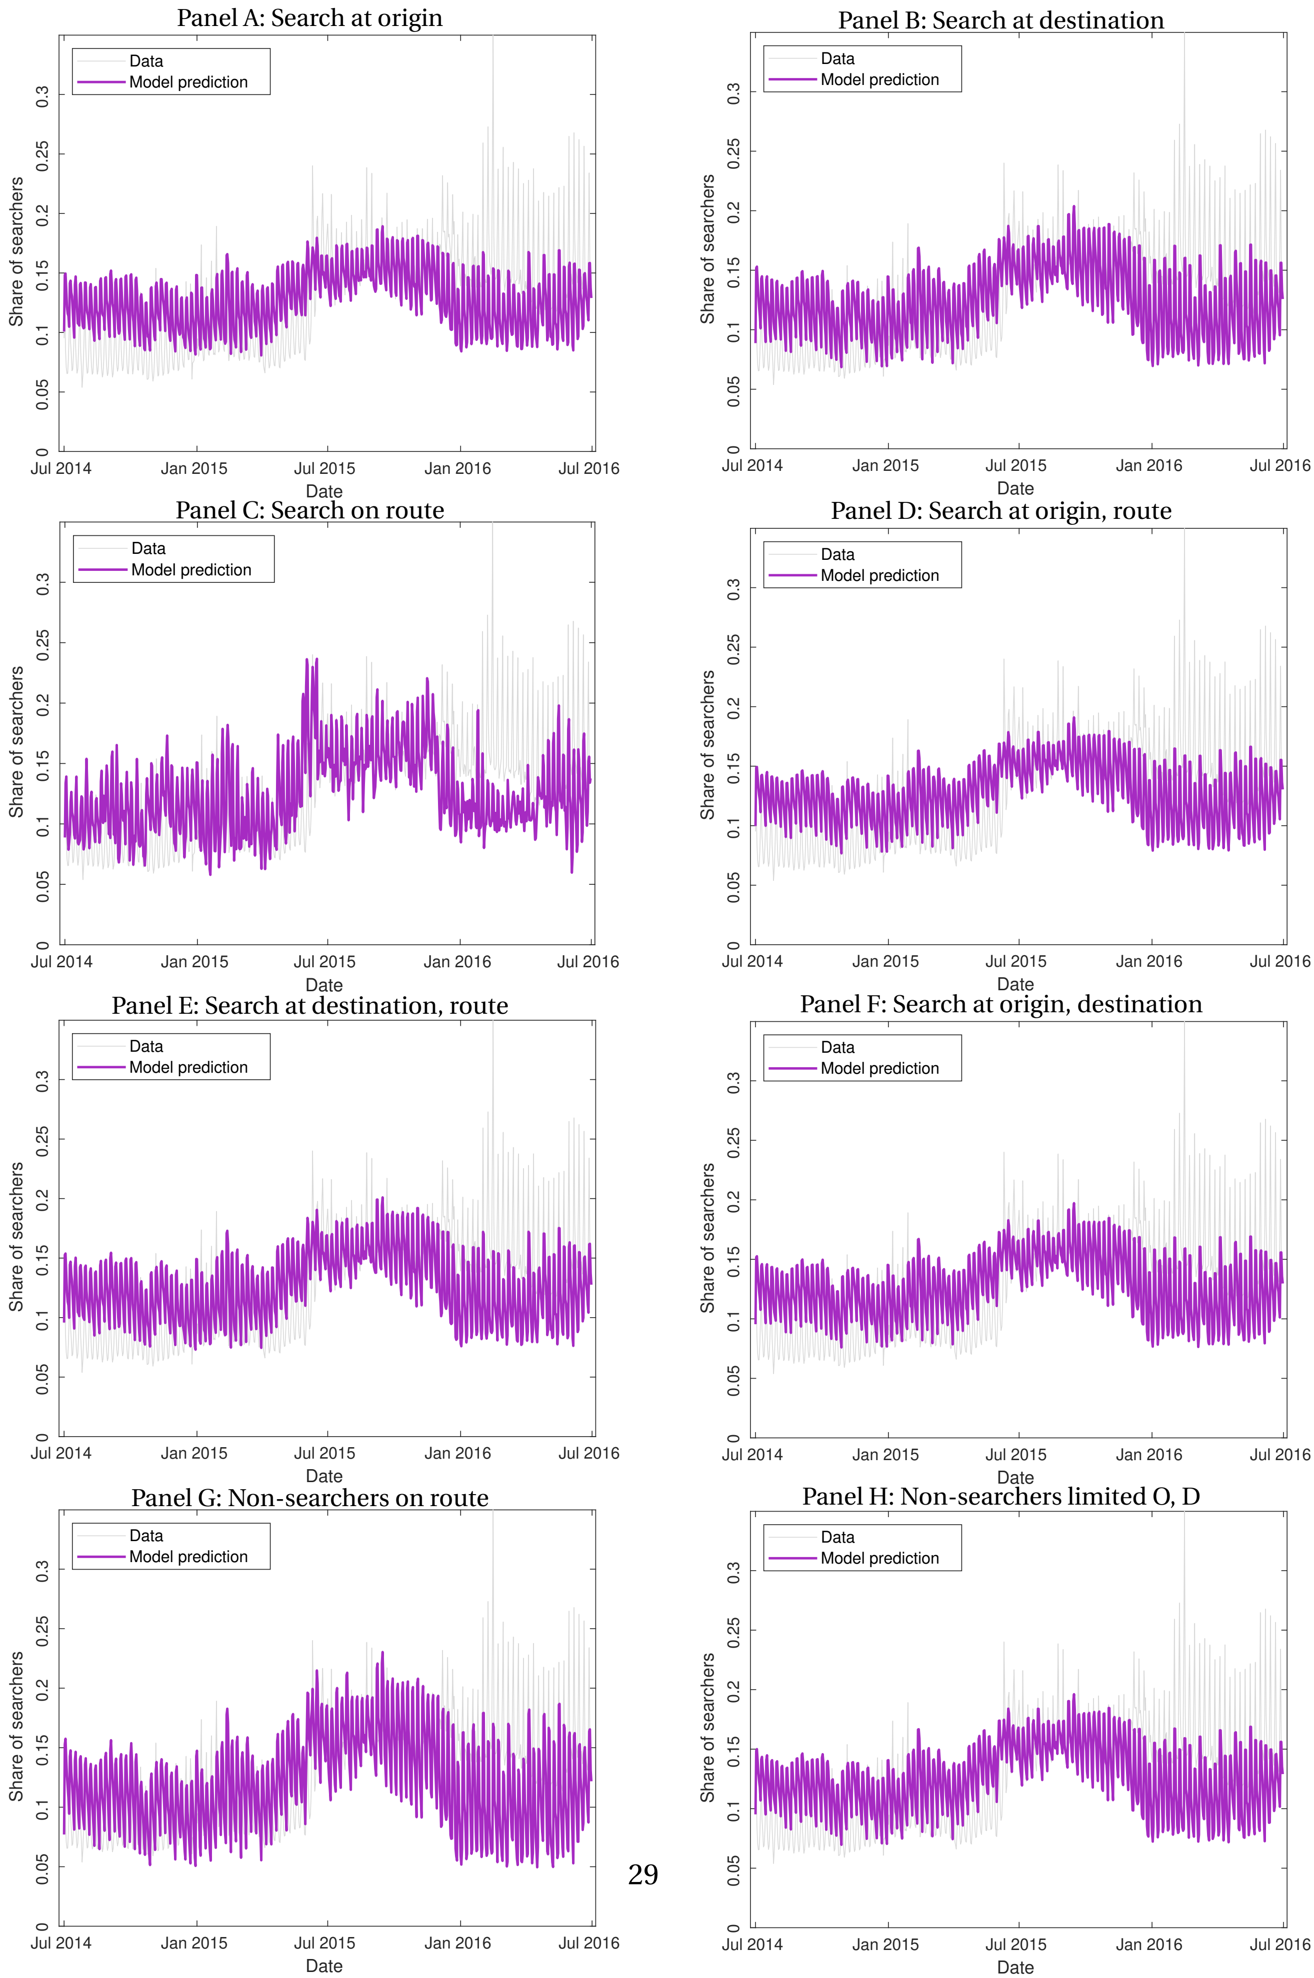

Figure D.7: Gains from search

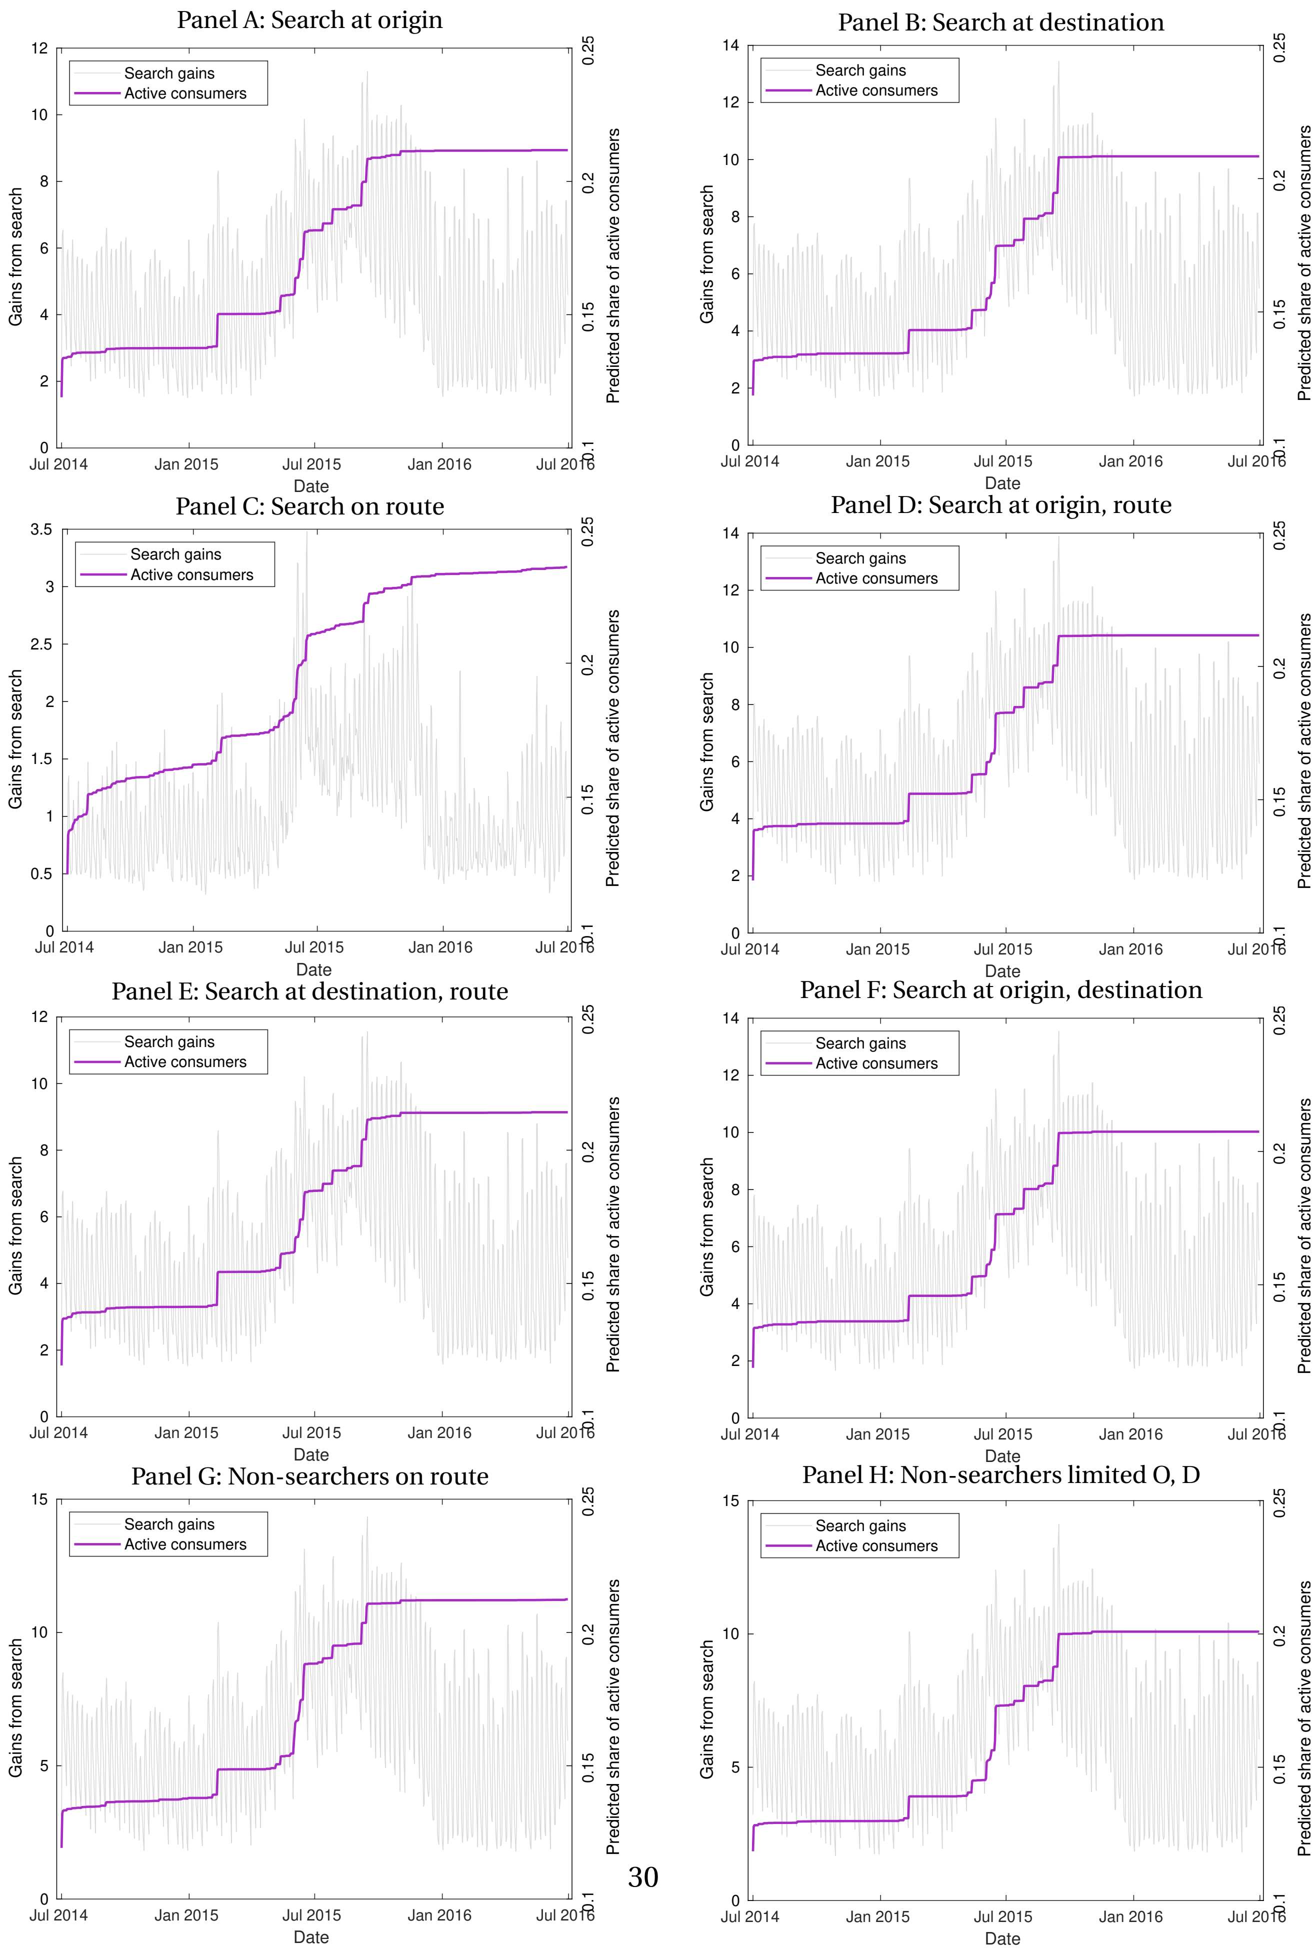

D.4 Alternative distribution assumptions

In this final set of model robustness checks, we estimate a sequence of variations of the model from the paper. Specifically, we use the Log Normal distribution instead of the Gamma distribution for startup and recurrent search costs, we allow for correlation between recurrent and startup costs, we allow for heterogeneity in the search radius employed by consumers, and we allow correlation between consumers' search costs and search radius.

For each consumer i, we draw a tuple of attributes, $\tau_i = (c_i, f_i, J_i, r_i)$. The recurrent and startup search costs c_i and f_i are drawn from a Log Normal distribution with mean μ_c and μ_f, and standard deviation σ_c and σ_f, respectively; the consideration set is drawn from the empirical distribution of origin-destination pairs from the 2016 Australian Bureau of Statistics O-D table; and the search radius r_i at the origin and destination is drawn from a truncated Log Normal distribution with mean μ_r and standard deviation σ_r. We discretize the search radius so that the radius is drawn from the set $\{1, 2, \ldots, 9, 10\}$ in kilometres.

Table D.3 contains estimates from this sequence of models. The first column of the table presents estimates where we use the Log Normal distribution in place of the Gamma distribution, and keep all other elements of the model unchanged, including assuming c_i and f_i are independently drawn. In column 2 ($c + f$), we introduce a correlation parameter, ρ, between the startup and recurrent search cost distributions. The model of column 3 incorporates heterogeneity in the size of consideration sets across consumers within a district. Thus, when drawing the consideration set J_i, we select both an origin-destination pair, and also the search radius r_i at the origin and destination which governs the size of J_i. In this specification, we assume c_i, f_i, and r_i are all independent draws. The models in columns 4 and 5 build on this model by introducing a correlation parameter with c_i and r_i (column 4, $c + r$) and f_i and r_i (column 5, $f + r$). For each specification, the top panel presents estimates from the full model including startup search costs, and the bottom panel contains estimates from the restricted model without startup search costs.

Estimation results are further illustrated in Figures D.8 - D.11. Figures D.8 and D.9 depict the startup and recurrent search cost distributions, respectively. Figures D.10 and D.11 illustrate the fit of the model for the full model and the restricted model omitting startup search costs, respectively. Each figure illustrates the baseline Log Normal model (Panel A), the Log Normal model with independent variation in search radius (Panel B), the Log Normal model allowing for correlation between search radius and recurrent search costs (Panel C), and the Log Normal model allowing for correlation between search radius and startup search costs (Panel D).

Table D.3: Estimation Results with Alternative Correlation Specifications

	Correlation Structure				
	None	$c+f$	None	$c+r$	$f+r$
	WITH STARTUP COSTS				
Recurrent costs					
μ_c	0.814	0.194	-0.779	0.777	-1.036
	(0.499)	(0.395)	(0.309)	(0.675)	(0.305)
σ_c	3.658	3.185	3.056	4.021	2.956
	(0.597)	(0.584)	(0.640)	(1.054)	(0.708)
Startup costs					
μ_f	3.257	3.195	3.008	2.952	4.007
	(0.055)	(0.088)	(0.090)	(0.190)	(0.503)
σ_f	0.706	0.615	1.038	1.038	1.007
	(0.059)	(0.079)	(0.064)	(0.072)	(0.093)
Search radius					
μ_r			-0.523	-0.259	-0.779
			(0.119)	(0.033)	(0.060)
σ_r			0.181	0.213	0.392
			(0.069)	(0.065)	(0.050)
Correlation					
ρ		-0.203		-0.192	-0.409
		(0.291)		(0.203)	(0.208)
Objective, $G(\hat{\theta})$	0.553	0.553	0.515	0.515	0.514
	NO STARTUP COSTS				
Recurrent costs					
μ_c	8.290		6.439	4.561	
	(0.543)		(0.499)	(0.985)	
σ_c	5.825		6.252	5.852	
	(0.485)		(0.531)	(0.603)	
Search radius					
μ_r			-0.776	-1.025	
			(0.176)	(0.111)	
σ_r			0.205	0.415	
			(0.187)	(0.082)	
Correlation					
ρ				0.207	
				(0.201)	
Objective, $G(\hat{\theta})$	1.219		0.985	0.988	

Notes: Robust standard errors are in parentheses (). The number of observations is $T = 731$ dates. All calculations assume consumers purchase 50 liters of gasoline. The log of recurrent costs, startup costs and search radius have distributions $N(\mu_c, \sigma_c^2)$, $N(\mu_f, \sigma_f^2)$, $N(\mu_r, \sigma_r^2)$.

Consider first the baseline Log Normal model, shown in column 1 of Table D.3 and Panel A of each figure. The results are very similar to those obtained using the Gamma distribution in the paper. The fit of the model is almost the same for both the full model and the model without startup costs. This can be seen by comparing the objective function (Table D.3 for the Log Normal model, and Table 2 for the Gamma distribution) and by comparing the model fit illustrated in Panel A of Figures D.10 and D.11 for the Log Normal model with that of the Gamma model in Panels A and B of Figure 7.

We allow for correlation between recurrent and startup search costs in column 2 of Table D.3. The correlation parameter is not statistically significant and this model adds no explanatory power beyond our baseline model. Relative to the baseline model, the recurrent search cost distribution has a lower estimated mean, and there is a mild negative correlation between startup and recurrent search costs. The fit of this model to the data is very similar to the baseline model, and we have not included it in the figures for the sake of brevity.

In column 3 of Table D.3 and Panel B of each figure, we present the model that incorporates heterogeneity in search radius across consumers. The estimated negative mean search radius parameter suggests that many consumers have the lowest possible search radius. This leads to smaller consideration sets relative to the baseline model, and therefore lower search benefits and greater heterogeneity in search benefits across consumers. To rationalize the lower search benefits, estimates of recurrent search costs are also lower, but there is little impact on estimated startup search costs. The positive and statistically significant variance parameter for the search radius indicates heterogeneity in search radius across consumers, and therefore adds to the heterogeneity across consumers in search benefits. Turning to Figure D.10, the main implication of these differences is greater high-frequency variation in predicted search activity, and this shows up as a slight improvement in the objective function.

Columns 4 and 5 of Table D.3 allow for correlation between search radius and recurrent and startup search costs, respectively. Both models suggest a negative relationship between search costs and the search radius, but provide no noticeable improvement in the fit of the model, as indicated by the value of the objective function and in Figure D.10. The negative correlation parameters imply that those consumers with higher search costs will engage in lower search for two reasons. In addition to the direct effect of higher search costs, these consumers tend to consider a smaller set of stations for search, and therefore they also have lower search benefits. The combined effect is greater heterogeneity in search activity across consumers. However, the effect is small, leading to no noticable difference in the fit of the model.

Finally, the bottom panel of Table D.3 and Figure D.11 illustrate estimates for the restricted model without startup search costs. The estimates in column 3 suggest that most consumers have the smallest possible search radius of 1km, but suggest heterogeneity in the search radius.

A reduced search radius implies lower average search benefits. Counterbalancing this effect, we estimate lower average recurrent search costs. Thus, the primary effect of introducing heterogeneity in the search radius appears to be increased volatility in search activity. This can be seen by comparing Panels A and B of Figure D.11, and leads to a noticable improvement in the objective function. Column 4 of Table D.3 and Panel C of Figure D.11 presents estimates that allow for correlation between search costs and the search radius. The estimated correlation coefficient is positive, but not statistically significant, and there is no appreciable impact on the fit of the model.

Figure D.8: Startup Search Cost Distributions

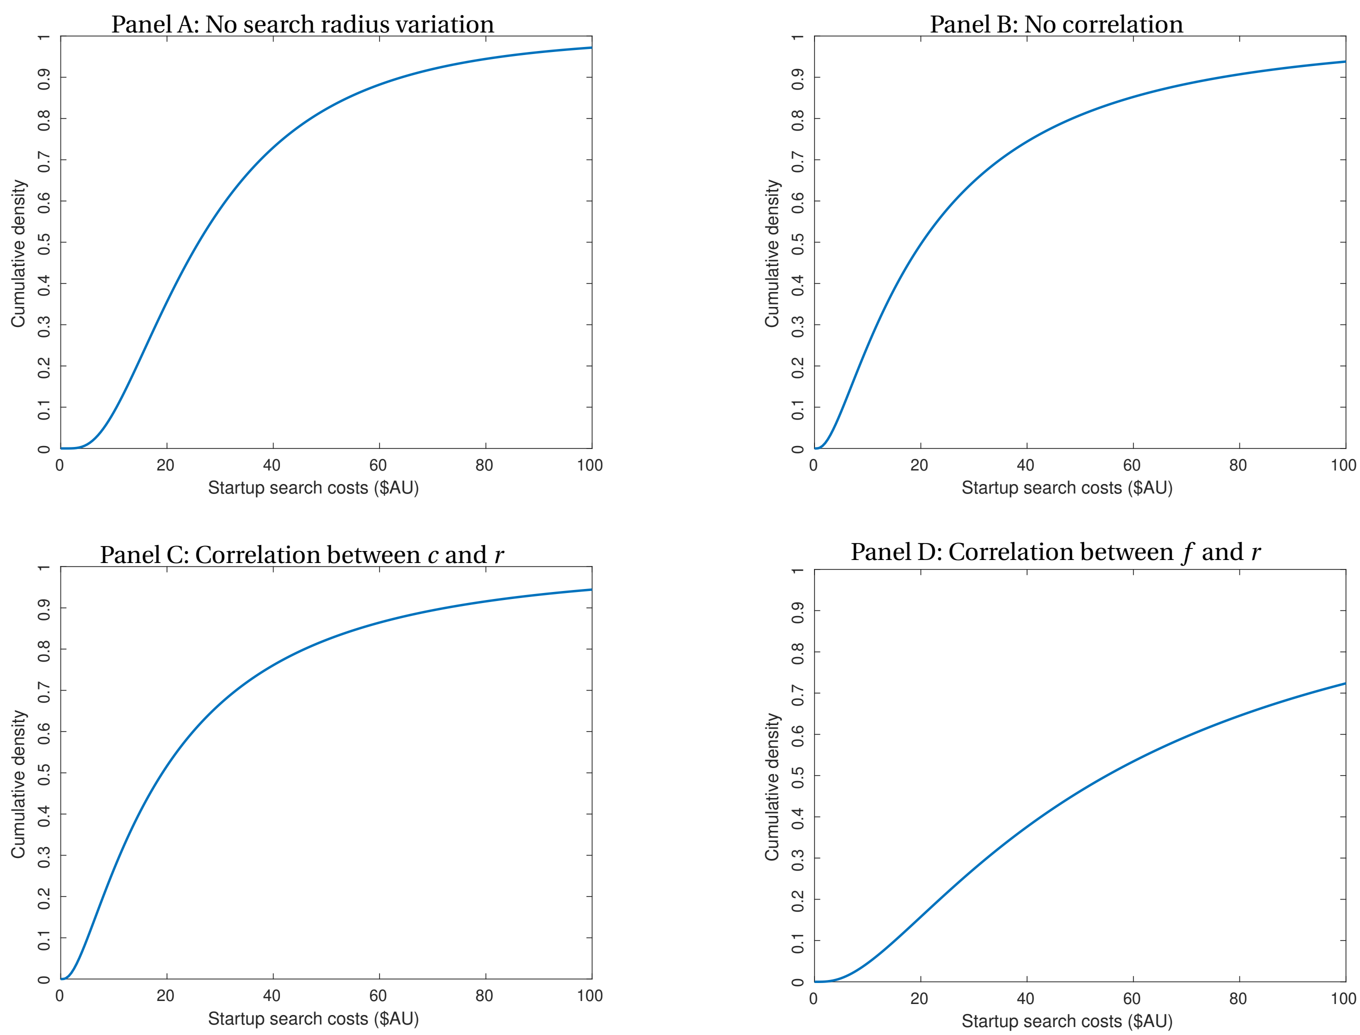

Figure D.9: Recurrent Search Cost Distributions

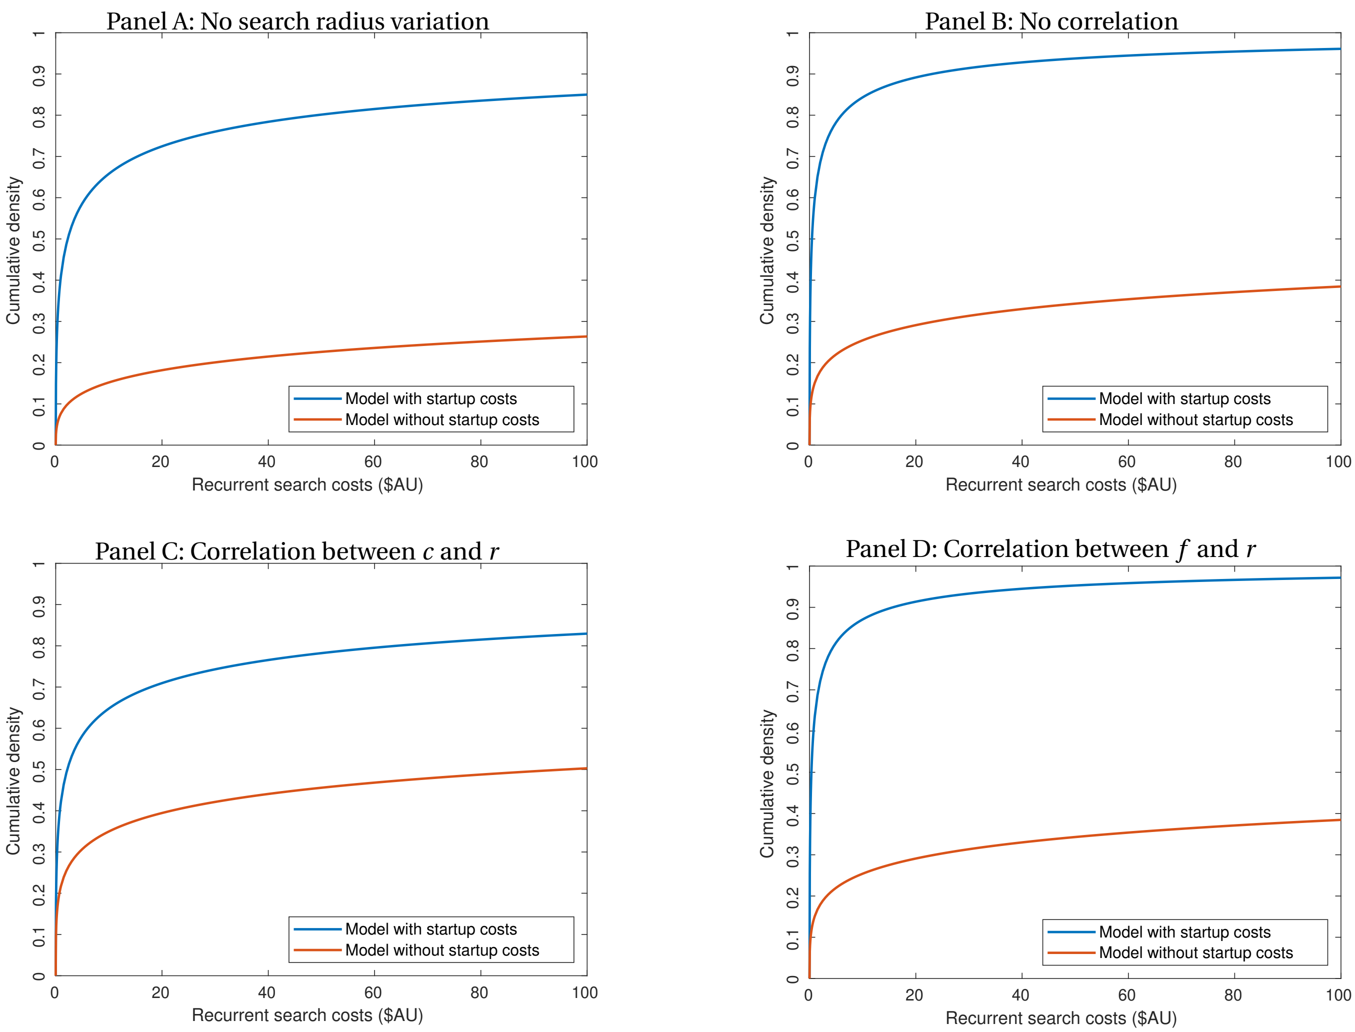

Figure D.10: Model Predictions, with Startup Costs

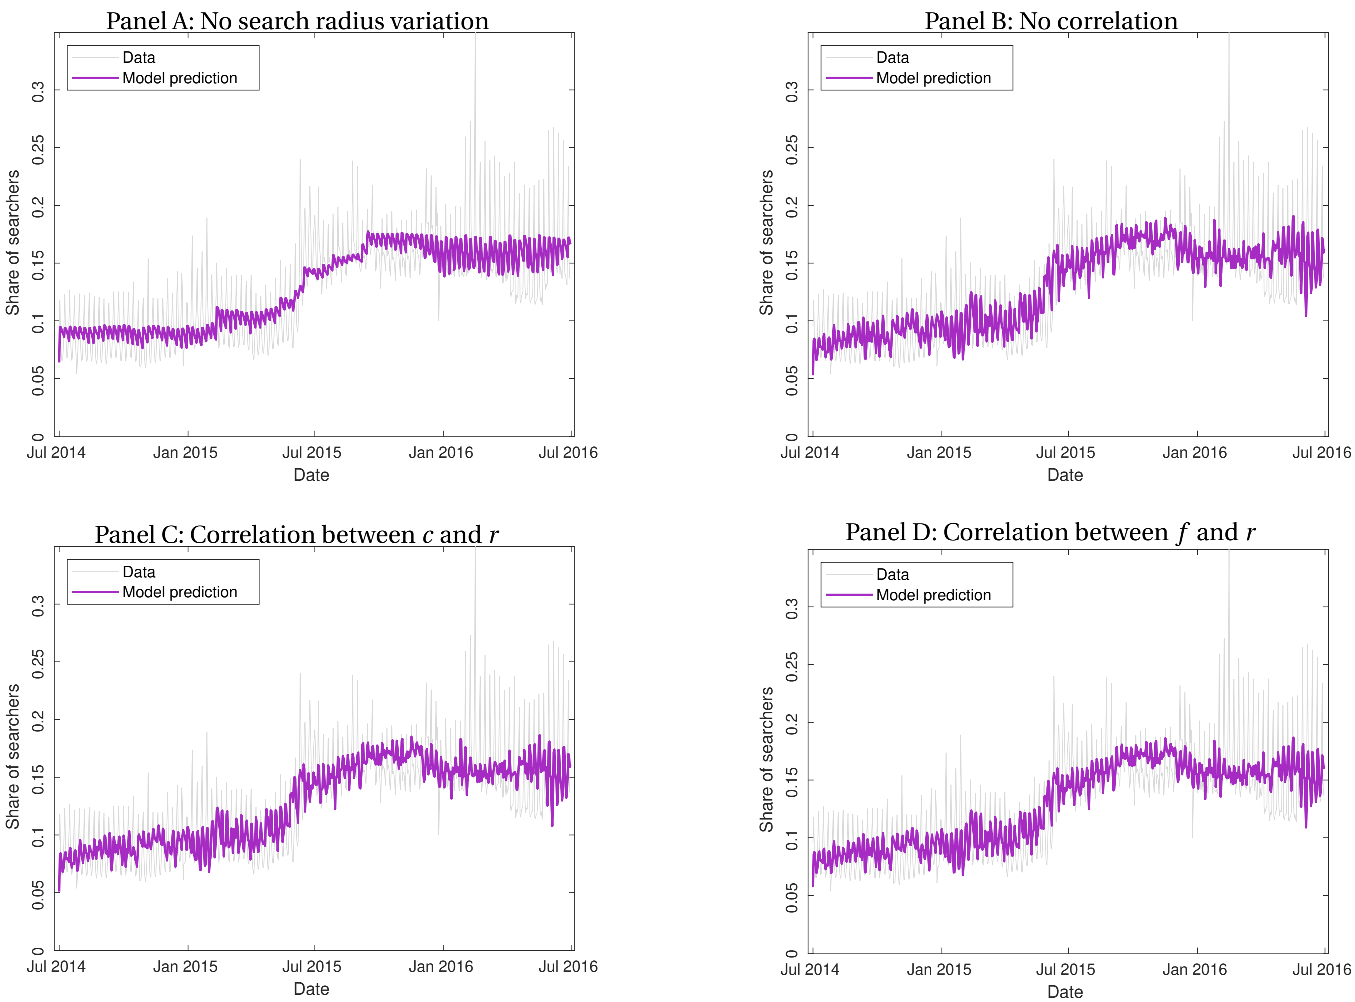

Figure D.11: Model Predictions, without Startup Costs

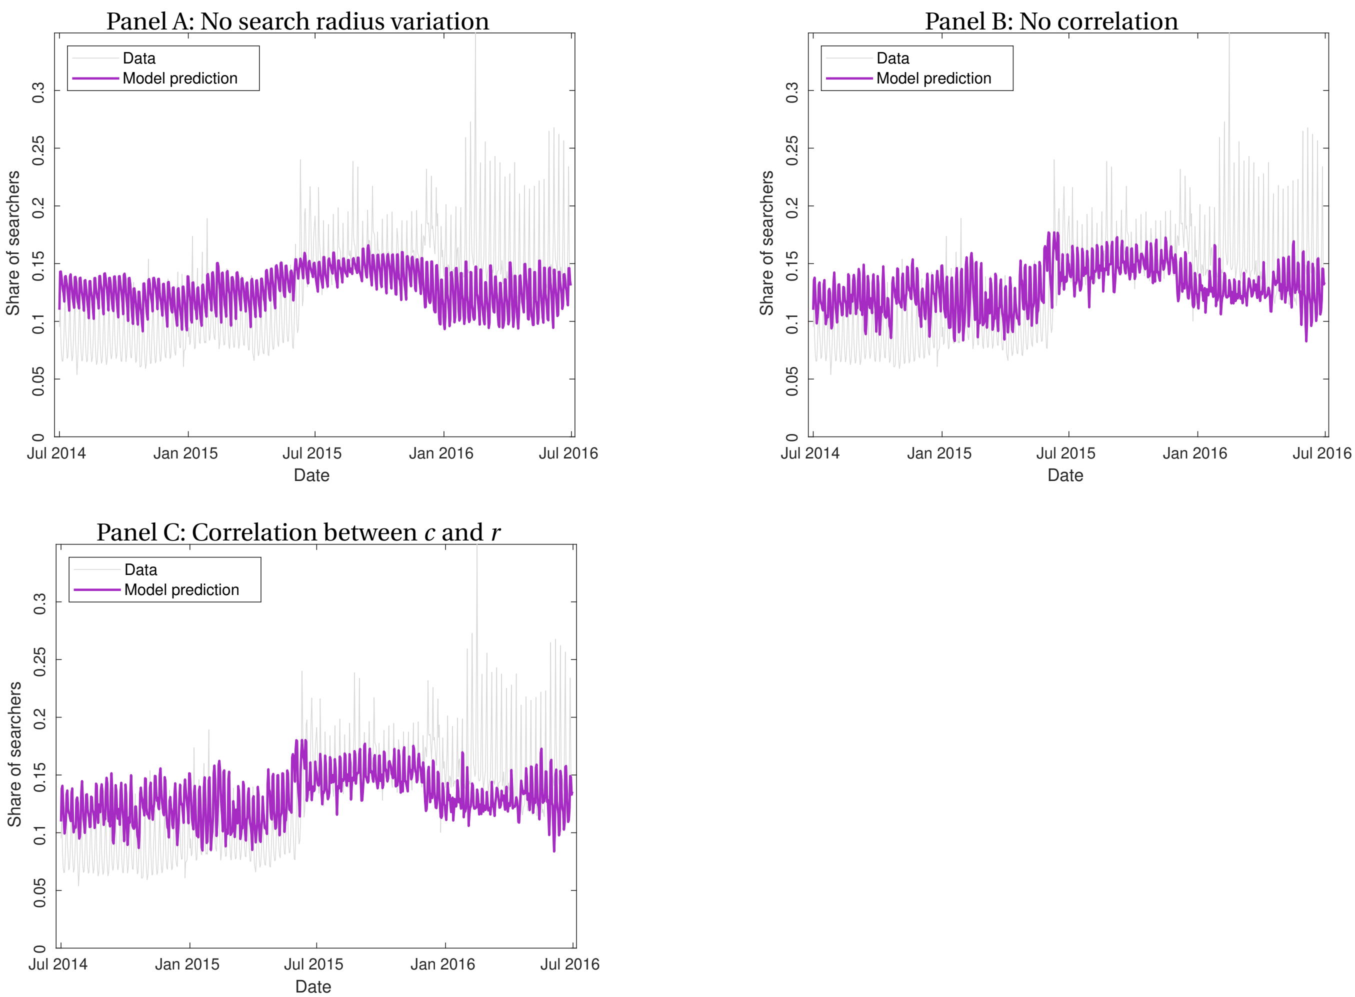

References

Donald W. K. Andrews. Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61:821–856, 1993.

David P. Byrne and Nicolas de Roos. Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109:591–619, 2019.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708, 1987.