

Search with Learning in the Retail Gasoline Market*

Xiaosong Wu[†]
Matthew S. Lewis[‡]
Frank A. Wolak[§]

November 2022

Abstract

This article estimates a model of optimal search where consumers learn the distribution of gasoline prices during their driving trips. Our model incorporates traffic information and leverages this ordered search environment to recover parameters of the search and learning process using only station-level price and market share data. We find that learning is a crucial component of search in this market. Consumers' prior beliefs regularly deviate from the true price distribution but are updated quickly following each new price observation. Counterfactuals reveal how these learning dynamics generate asymmetric search patterns commonly associated with asymmetric cost pass-through.

Keywords: consumer search, consumer learning, dynamic discrete choice model, asymmetric price adjustment, gasoline

JEL Classification: D83, L10, L13, L9

*We are grateful to the editor, Ying Fan, and three anonymous referees for comments and suggestions that significantly improved the article. We thank Babur De los Santos, Christy Zhou, Sergei Koulayev, Matthijs Wildenbeest, Aaron Barkley, and David Byrne for their valuable suggestions. We also thank seminar participants at Clemson University, the University of Melbourne, the 2019 Consumer search and switching cost workshop in Los Angeles, and the 2019 Asia-Pacific Industrial Organization Society in Tokyo for helpful comments. All errors are our own.

[†]University of Melbourne; andy.wu1@unimelb.edu.au.

[‡]Clemson University; mslewis@clemson.edu.

[§]Stanford University; wolak@zia.stanford.edu.

1 Introduction

Since Stigler (1961) and McCall (1970), consumer search models have played an important role in explaining imperfectly competitive behavior in many markets. Consumers in these models tradeoff the cost of searching to acquire additional price information against the expected benefit of search, derived from consumers' beliefs about the price distribution. Standard search models rely on the convenient, yet strong, assumption that consumers *know* the true price distribution, thereby simplifying the calculation of consumers' gains from search. In many cases, however, when the prices in a market are unfamiliar to consumers or the price distribution changes regularly with market conditions, consumers are unlikely to know the price distribution with any degree of certainty.¹ Rothschild (1974), Dana (1994), and Benabou and Gertner (1993), among others, have relaxed this assumption, developing theoretical models of search with learning where consumers engage in costly search not only to reveal the prices of particular sellers but also to learn about the actual price distribution. Nevertheless, empirical studies of search behavior have largely continued to leverage the assumption that consumers search from a known price distribution.

In this article, we relax the assumption of a known price distribution and estimate a model of optimal search by consumers who may be unaware of the true price distribution but update their prior beliefs as they search. We are able to estimate the parameters governing the consumer learning process by taking advantage of the fact that price draws occur in a known sequence. In many contexts, prices are revealed based on the order in which consumers encounter different sellers as they navigate through the marketplace. For example, consumers observe prices in a specific order as they pass sellers within a market or scroll down a list of products on an online shopping website. In our setting, we leverage a crucial observation in the retail gasoline market: consumers are likely to search and learn the distribution of prices during their driving trips. This feature allows us to recover the parameters of our model from observed prices, station market shares, and the volume of traffic that passes each gas station.

Estimates of this model offer several new insights. First, in contrast to the known price

¹For example, Matsumoto and Spence (2016) and Jindal and Aribarg (2021) use survey and experiments to elicit price beliefs and find that consumers have prior price beliefs different from the actual price distribution and update their beliefs in response to search outcomes.

distribution assumption, we find that consumers' initial priors often differ significantly from the true price distribution, resulting in what we refer to as a *biased prior*.² Second, we find that consumers are relatively uncertain about their prior beliefs and, therefore, learn quickly from observed prices. Third, the model reveals how biased beliefs and learning influence search behavior and demand. These insights clarify a mechanism through which price fluctuations can asymmetrically influence search. Such asymmetric search patterns are commonly cited as an explanation for why cost increases and decreases are passed through asymmetrically in a wide variety of product markets (Peltzman, 2000).³

The retail gasoline market is an ideal environment to study consumer search and learning behavior. Frequent price changes resulting from a volatile wholesale cost, as presented in Figure 1, make it difficult for consumers to maintain accurate information on each station's price as well as the distribution of these prices in the market. Our analysis of search introduces two important components that are likely to characterize consumers in this environment. First, consumers are assumed to be uncertain about the price distribution. Second, consumers' prior beliefs are likely to differ from the empirical price distribution.

To capture these features, we propose a sequential search model with learning that builds on Rothschild (1974), emphasizing spatial and ex-ante vertical differentiation of sellers. Forward-looking consumers start from diffuse prior beliefs likely influenced by prices observed during past driving trips. As consumers encounter a new price observation along their predetermined travel route, they update their beliefs about the price distribution in a Bayesian fashion before deciding whether to purchase gasoline or continue searching. Consumers stop at a station when the realized utility is higher than the continuation value of search conditional on their posterior beliefs. The continuation value of search summarizes the expected value of purchasing at the remaining stations along a route and the alternative of waiting to purchase during a future trip where they might encounter better offers. However, postponing a purchase becomes difficult if one is low on gas. Thus, the search friction in this market takes the form of postponement costs.

²For brevity, we use prior bias to refer to the notion that consumers' prior mean of the price level does not necessarily match the true average price. It is rational for consumers to have beliefs different from the true price distribution when they only observe past prices and noisy signals of current prices.

³Asymmetric pass-through is particularly common and well documented in gasoline markets (Borenstein et al., 1997; Lewis, 2011) and is similarly prevalent in our sample (see Figure 1).

To estimate the model, we utilize a panel dataset of station-specific prices and quantities for gasoline stations in a small city from December 20, 2014, to May 31, 2016, and combine it with data on traffic flows in the city. Based on the assumption that price search is ordered and determined by driving patterns, we can construct an empirical distribution of search sequences of gasoline stations using the traffic data. Total daily gasoline sales at each station are then modeled as an aggregation of the purchase decisions of individuals searching and learning along different travel routes.

Our novel utilization of traffic data also allows us to more realistically model the search behaviors in the retail gasoline market without losing tractability. By replacing the common random sampling assumption with ordered search determined by observed traffic flows, our model allows substitution patterns to depend on the amount of traffic stations share. In addition, we are able to introduce publicly observable vertical differentiation of sellers, allowing consumers to be familiar with the time-invariant characteristics of the stations they regularly encounter along their driving routes.

Our search with learning model nests the standard search model with a known price distribution, providing us the opportunity to test empirically the assumption of a known price distribution in the context of the retail gasoline market. Our estimation results suggest that consumers' initial prior beliefs are significantly biased. Specifically, the average absolute difference between the estimated prior mean and the actual price level is 2.7 cents per gallon (cpg), approximately 3.3 times the size of the average day-to-day price change. However, consumers put relatively little weight on these priors, updating beliefs rather quickly and considerably reducing the bias after a few current price observations. The findings overwhelmingly reject the null hypothesis of a known price distribution which assumes both a correct price belief and no learning. They also highlight the importance of accounting for learning when analyzing search behavior. Learning occurs rapidly in our context, and models with learning produce much more accurate predictions of consumer behavior. Estimates from a restricted version of the model with no learning fail to identify prior bias and overestimate the median postponement costs by approximately 33 percent.

Estimating a structural model of search with learning also provides a powerful framework for examining the nature of spatial competition in the market. Demand is estimated

to be highly elastic at the station level. A typical station has an estimated own-price elasticity of -8, similar to the findings in Wang (2009). In addition, the traffic data allow the model to generate realistic substitution patterns across stations. If an additional 15 percent of a station’s passing traffic has previously driven past a neighbor station, the cross-price elasticity between the two stations is predicted to be 0.64 higher, sufficient to move the station pair from the 5th percentile to the 95th percentile of the cross-price elasticity distribution.

Incorporating learning into the search environment helps explain important patterns that alternative models can not capture. For example, a non-trivial share of competing stations is found to have negative cross-price elasticities of demand, which can arise because the information conveyed by one station’s price can impact consumers’ beliefs about prices at subsequent stations. Learning also creates an environment where station-level demand elasticities can change over time with price fluctuations. We estimate demand to be more elastic when prices are rising than when prices are falling, creating an environment in which cost changes can be passed through asymmetrically. In a counterfactual exercise, we investigate how the learning process influences demand asymmetry. We find that when past price levels more heavily influence prior beliefs about the distribution of prices, demand elasticities respond more asymmetrically to price changes. However, conditional on the degree of prior bias, higher prior uncertainty leads to less asymmetric demand.

The remainder of the article is organized as follows. The next section places our work in the context of related literature. Section 3 introduces the data used to estimate our model. Section 4 presents descriptive statistics and key features of the market that motivate our model. The model of search with learning is introduced in Section 5, and the estimation strategy and model identification process are discussed in Section 6. Section 7 presents the estimation results. Section 8 discusses our counterfactual analysis. The last section concludes.

2 Related Literature

Much of the empirical literature on consumer search quantifies search frictions and emphasizes their importance in various markets based on the assumption that consumers know

the distribution of offers or match values (e.g., Hortaçsu and Syverson, 2004; Hong and Shum, 2006; De los Santos et al., 2012; Koulayev, 2014; Honka, 2014; Nishida and Remer, 2018; Lin and Wildenbeest, 2020; Moraga-González et al., forthcoming).⁴ We build on this literature by incorporating consumers that learn about the true price distribution as they search, and show this to be an important aspect of behavior in the retail gasoline market. Our work adds to a developing body of empirical research on consumer search with learning. Both Koulayev (2013) and De los Santos et al. (2017) empirically analyze models of search with learning and show that ignoring learning can bias search cost and elasticity estimates. However, they do not estimate the learning process and take prior beliefs as given. In contrast, we develop an empirical strategy to identify both prior uncertainty and prior bias using only aggregate data. Several more recent studies have modeled learning behavior in settings different from ours. Ursu et al. (2020) estimate a sequential search model where consumers search to learn their individual match values for restaurants on a review website. Their estimates suggest a high prior uncertainty that rationalizes the considerable time consumers spent searching each restaurant. Hu et al. (2019) develop a dynamic model of search and Dirichlet learning to study consumers’ purchase behavior on Groupon. They find that new consumers have an overly optimistic prior about the distribution of deal quality. Through their interaction with the website over time, consumers have more certain and accurate beliefs about the quality distribution. Consumer learning explains the observed declines in click-throughs and increases in conditional purchase probability. However, unlike our study, estimation in each of these articles requires individual search and purchase history data.

The underlying consumer search process and identification method employed in our model also differ from the existing literature. A number of studies have developed methodologies to estimate search costs using only aggregate data (e.g., Hortaçsu and Syverson 2004; Hong and Shum 2006; Moraga-González and Wildenbeest 2008; Wildenbeest 2011). These studies overcome the curse of dimensionality when integrating over the unobserved search sequences by applying the assumption of random sampling and ex-ante product homogeneity. We propose a new estimation strategy for settings where the order of price

⁴See Ellison (2016) and Honka et al. (2019) for a review of the studies on consumer search.

observations is determined by how consumers navigate through the marketplace. In particular, we replace the random sampling assumption with variation in search sequences identified using data on driving patterns. With this approach, we can introduce learning and ex-ante seller differentiation into a sequential search setting without losing tractability. Additionally, we can allow for more realistic substitution patterns between stations. Our search technology also differs from the literature on sequential search for ex-ante differentiated products (e.g., Weitzman 1979; Kim et al. 2010, 2017) in that we model consumer search order as exogenously given by the traffic data rather than endogenously determined by the decreasing order of reservation utilities.⁵

Moreover, our article provides valuable insights into extensive literature on retail gasoline price dynamics. A large body of empirical research provides evidence of asymmetric cost pass-through in the retail gasoline market (e.g., Borenstein et al. 1997; Lewis and Noel 2011; Byrne 2019). Tappata (2009), Yang and Ye (2008), and Lewis (2011) develop theoretical models showing such pricing behavior can arise when consumers have imperfect knowledge of the price distribution. Lewis and Marvel (2011) and Byrne and de Roos (2017) offer evidence of the influence of price movements on search activity and illustrate that observed patterns of asymmetric cost pass-through and fluctuations in price dispersion are consistent with the search dynamics. However, little is known about consumers' price beliefs. Our structural model contributes to this literature by estimating how consumers form their price beliefs in this market. We demonstrate how learning primitives, prior bias and prior uncertainty, cause the intensity of consumer search and the elasticity of demand faced by stations to change over time in response to gas price fluctuations. Therefore, this article brings together two streams of literature, structural analysis of consumer search and research on cost pass-through in the retail gasoline market.

Our analysis also relates to the broad set of studies examining and modeling spatial competition and its consequences, particularly those using spatial information on consumers to identify the intensity of competition (e.g., Smith 2004; Thomadsen 2005; Davis 2006; Manuszak and Moul 2009; Houde 2012; Miller and Osborne 2014). These studies incorporate the distance between sellers and consumers into a discrete-choice demand framework

⁵Our ordered search can be interpreted as a special case of Weitzman's sequential search where the costs of deviating from the current travel route are much larger than the potential gains.

while assuming full information. In particular, our model is most similar to that of Houde (2012), who also uses road network structure and traffic flow volume to determine the degree of spatial competition between stations. However, we incorporate imperfect price information and learning, allowing consumers' expectations to be influenced by past price levels. Therefore, a unique feature of our model is that a station's demand and the elasticity it faces can change over time with fluctuations in consumers' beliefs about the price distribution. Moreover, search and learning allow for a different structure of cross-price elasticities than the full-information spatial model of Houde (2012). For example, in some circumstances, negative cross-price elasticities can arise between competing sellers as a result of the information conveyed through prices.

3 Data

We use aggregate data to make inferences about consumers' search and learning behavior that leads to gasoline purchases. Our sample consists of 46 gasoline stations in a small city with an urbanized area population of approximately 75,000.⁶ The sample period runs from December 20, 2014, to May 31, 2016, for a total of 529 days, during which time we observe the daily price of gasoline at all 46 stations and the daily gasoline transaction volume for 33 of these stations.⁷

We complement the gasoline price and quantity data with data on vehicle traffic flows for our sample region. We use this traffic data to construct an search sequences of stations in the city. The following subsections describe the three primary data sources used for our empirical analysis.

Gasoline Price Data

The per-gallon price of regular unleaded gasoline is collected from two separate gasoline price reporting websites. The primary source is MapQuest.com, an online web mapping ser-

⁶The city name is not disclosed to protect the identities of the gas stations.

⁷Because our primary focus is to study search behavior for gasoline, we exclude 14 mom-and-pop establishments that operate primarily as convenience stores and have a gasoline sales volume lower than the smallest station in the sample.

vice whose gasoline price data are provided by the Oil Price Information Service (OPIS).⁸ We record prices from MapQuest.com once per day for every station in the city. Unfortunately, MapQuest (OPIS) does not update every station’s price daily. On average, a station’s price is updated on 54 percent of the sampled days.⁹ To address the issue, we complement MapQuest.com’s data with price data collected daily from GasBuddy.com.¹⁰ Unlike MapQuest.com, prices on GasBuddy.com are reported by volunteer spotters in the area. To minimize any issues caused by the potential inaccuracy of prices reported on GasBuddy.com, we only use prices from GasBuddy.com when MapQuest.com does not report the corresponding price for that station on that day.¹¹ Stations are matched across the two data sources based on the geographic coordinates of the stations, cross-validated with Google Map’s geographic coordinates to ensure accuracy.¹² After merging the price data from these two sources, station prices are missing for only 9.2 percent of the sample days. The remaining missing prices are replaced with the most recent price observed at that station. The average duration over which prices are imputed is 1.6 days.¹³ Besides price data, we also obtain information on station characteristics from these sources, including name, brand, address, and geographic coordinates. Moreover, we visit Google Street View and manually collect additional information such as the number of islands and pumps and street conditions for each station.

⁸OPIS obtains price information from credit card transactions and direct feeds from gas stations.

⁹The price coverage rate is slightly lower than other studies that use OPIS data. A possible reason is that the sample city is mid-sized and has more low-volume stations than the major cities studied by other researchers. Fewer credit card transactions result in fewer price feeds to OPIS.

¹⁰Gasoline price data collected from MapQuest.com and GasBuddy.com are widely used in the literature on retail gasoline prices, for example, Lewis and Marvel (2011) and Remer (2015).

¹¹Atkinson (2008) shows that prices on GasBuddy.com can accurately identify the features of retail gasoline price competition despite occasional errors. GasBuddy.com price data match that from MapQuest.com for 76 percent of the days when both are available. A closer data investigation reveals that most unmatched prices are likely due to intra-day price changes.

¹²A station’s name or address cannot be used as a unique identifier for the matching because a station’s name is not unique to a station, and different websites may use different aliases for a street or highway.

¹³Our estimation cannot accommodate missing prices because a station’s price affects many stations’ sales through the traffic network. One missing price will result in a large number of lost observations.

Gasoline Transaction Data

Daily station-level expenditure data have been obtained from a major financial services provider for 33 of the 46 stations in our price sample.¹⁴ These data reflect the total dollar amount of purchases made using debit and credit cards associated with the provider’s purchase processing network at each station on each day. Pay-at-pump and in-store purchase totals are reported separately. To eliminate the measurement error caused by non-gasoline transactions, we use pay-at-pump transactions only. A daily measure of the quantity of gasoline purchased at each station is constructed by dividing the total pay-at-pump expenditures by the price of regular unleaded gasoline at the station on that day.¹⁵ Although this quantity measure excludes gasoline purchased with cash or in the store, around 72 percent of consumers purchase gasoline at the pump (NACS 2016 Retail Fuels Report). Therefore, we believe that our measure of the quantity of gasoline transacted at each station reasonably describes the behavior of consumers searching for and purchasing gasoline.

Empirical Distribution of Search Routes

As individuals drive along their travel routes, the decision to purchase gasoline at a particular station is affected by the prices observed up to this station as well as the characteristics of the remaining stations along the route. Consequently, we model consumers’ search and purchase decisions at the search-route level. A search route is defined as a unique ordered sequence of stations visited, exogenously determined by consumers’ travel needs.

The empirical distribution of search routes describes the predicted share of drivers traveling along each possible ordered sequence of stations on an average day. Its construction involves two elements: (i) the number of drivers traveling from an origin to a destination and (ii) the route drivers take along the street network connecting the two points.

For the first element, we use the origin-destination travel demand estimates for local residents produced by the state Department of Transportation,¹⁶ which report an estimate

¹⁴The name of the provider as well as the station names and locations in the data are withheld to protect confidentiality.

¹⁵This construction introduces potential measurement error, as it overestimates the quantity transacted for mid-grade and premium gasoline, which have higher prices. However, it has been estimated that only 15 percent of gasoline transactions are mid-grade or premium.

¹⁶The Origin and Destination Table is an output of the travel demand model constructed by the De-

of the average number of drivers traveling from one Traffic Analysis Zone (TAZ) to another TAZ. The origin-destination table spans a seven-county area around the focus city and contains approximately 1800 TAZs. Most TAZs are relatively small, with 75 percent of the traffic zones occupying an area of less than 1.5 km^2 . The larger traffic zones have few residents and are at the fringe of the counties.

To compute the route drivers take traveling from an origin TAZ to a destination TAZ, we assume that all drivers take the route that minimizes driving time. We select the centroid of the origin and destination TAZ as the drivers' start and end locations and determine the single fastest travel route for every origin and destination TAZ pair based on the street network in the area.¹⁷ The ArcGIS Network Analyst package is used to calculate the fastest travel route, with road network data obtained from ArcGIS StreetMap North America.¹⁸

Next, we identify the stations along each travel route and the order in which they are passed. Figure 2 provides an example of a travel route connecting a starting location A and an ending location B, including the three stations available to drivers on this route. In many cases, it is difficult for drivers to visit stations on the opposite side of the street because some left-turns cannot be easily made. To consider the potential cost of making left-turns in the model, we also record the side of the street a station is on along each route. We discuss the different left-turn types and their difficulties in the next section.

A search route in our model is formally defined by a specific sequence of stations. As multiple travel routes (origin and destination pairs) may pass the same set of stations in the same order, travel routes are aggregated to the search-route level.¹⁹ A total of 991 search routes are identified at the beginning of the sample, and the number increases to 1046 after two additional stations enter the market.²⁰ The number of travelers on a search route is constructed by summing up all drivers traveling past the same ordered set of stations (and

partment of Transportation to forecast the traffic in the year 2020.

¹⁷To reduce the computation burden, we grouped the TAZs in each of the surrounding counties into 8 clusters of TAZs based on their locations using the K-Mean algorithm.

¹⁸ArcGIS Network Analyst extension: <https://www.esri.com/en-us/arcgis/products/arcgis-network-analyst/overview>. StreetMap North America: <https://www.arcgis.com/home/group.html?id=ddd06a0bde9c45a1b3e786a2b4e695e8#overview>.

¹⁹Travel routes passing no stations are excluded from the sample. We also exclude search routes with fewer than 20 daily drivers.

²⁰In the structural estimation, the sample periods are divided into three parts based on the entry date. The empirical analysis is based on the empirical distribution of search routes in each period, respectively.

only those stations). Dividing it by the total number of daily travelers in the area produces a vector of the share of travelers on each search route.

Although our daily station-level gasoline expenditure data provide many advantages, there are a few limitations. First, expenditure data are only observed for around 70 percent of the stations in the market. Our model predicts demand at every station, but the identification is based on the stations with observed quantity data. Variation in observed station characteristics is also limited, so the vertical differentiation of sellers is incorporated into the model based on station type. To maintain group size and protect each retailer's identity, we apply three brand dummies to account for brand heterogeneity: one dummy for major-branded stations and two for retailer brands. The remaining stations are collectively classified as generic stations. Additionally, we include a small station dummy and a large-format station dummy to control for the station scale.

Stations located near Interstate Highway exits also present a challenge. The origin-destination traffic data only describe the travel patterns of local drivers, so potential demand from Interstate drivers is not accurately reflected. Interstate drivers also observe prices and make purchase decisions very differently than the local drivers modeled in this article. To more accurately capture demand at these stations, our model includes a separate dummy variable to account for the average differences between the model predicted and the observed market shares for each station located at an Interstate Highway exit.

4 Retail Gasoline Market Overview

Before introducing the structural model, we discuss the features in the retail gasoline market that motivate our modeling choice. More specifically, we first examine the relationship between the station average transaction volume and station characteristics. We then discuss why consumers are likely to have imperfect price information and why it is important to incorporate learning when modeling consumer search in this market.

Station Transaction Volume and Station Characteristics

Table 1 summarizes the station-level average prices and quantities as well as some important station characteristics. The top panel provides statistics for all of the stations in the city, whereas the lower panels separately consider specific station types. Average gasoline prices vary somewhat across stations, exhibiting an interquartile range of 10 cents per gallon around a city-wide average of \$1.61 per gallon (before taxes). Considerably more heterogeneity is exhibited in station-level average transaction volume. Among the 33 stations for which we observe quantity data, the 75th percentile station sells 6.5 times more gasoline than the 25th percentile station. Major-branded stations such as Shell, BP, and Exxon, among others, account for approximately 37 percent of the stations in the city.

On some streets, it may be difficult for drivers traveling in a certain direction to visit stations on the opposite side of the street. For this reason, we classify three types of stations based on left-turn difficulty. Approximately 26 percent of our sample stations can be easily visited by drivers traveling on both sides of the street. These include stations on two-lane or multi-lane roads with a left-turn zone in the center. Another 28 percent of the stations are located where no left-turns are possible because the street has a physical curb or median in the center. The remaining stations are located at major intersections with a traffic light. Casual observation suggests that drivers are likely to forgo possible price savings at these stations to avoid waiting for the left-turn traffic light, especially when the intersection is busy. To provide a conservative measure of the number of consumers each station faces, we define a station's *direct traffic* as the number of drivers who can easily visit the given station, which includes drivers driving on the same side of the street as the station or on the opposite side of the street where a left-turn can be easily made without involving a traffic light. As shown by the last row of the top panel, 11.5 thousand drivers directly drive past a station on average.

Panel (a) of Figure 3 depicts a positive relationship between the direct traffic volume and the transaction volume at a station, both measured in logarithms,²¹ revealing that stations passed by more drivers also sell more gasoline.²² There is significant variation

²¹We exclude the stations at the exit of the interstate highway from this figure because we do not have data measuring highway traffic volume.

²²The correlation between the average gasoline transaction volume and the passing traffic volume is 0.28,

around this relationship, suggesting that other station attributes such as price reputation and brand quality may also influence station sales. Nevertheless, the pattern demonstrates the advantage of using traffic data to simulate consumers' search patterns for gas stations. Other empirical studies of consumer search (e.g., Hong and Shum 2006; Wildenbeest 2011; Nishida and Remer 2018) have typically adopted an equal-probability random sampling assumption when individual search histories and quantity data are not observed. Our data suggest that consumer search along travel routes better represents consumer behavior than the random sampling assumption.

In recent years, stations with a large number of islands and a large convenience store attached are becoming increasingly popular (e.g., Noel 2016). We group stations into three categories based on their scale: large-format retailers, small-sized stations, and mid-sized stations. In particular, we define large-format stations as retail stations with at least six islands.²³ All large-format stations in our sample have a sizable convenient store attached. In contrast, small stations have no more than three islands, with a small booth in the center. The remaining stations are categorized as mid-sized stations. The bottom two panels in Table 1 describe the price and quantity distributions for the small and large-format stations. Whereas the average price at large-format stations is, on average, six cpg cheaper than at small-sized stations, the average daily sales volume at large stations is 7.8 times greater than at small stations. Notably, large-format stations all have average prices in the lowest quartile of the city distribution, whereas their average sales volumes are all in the highest quartile. The negative correlation between stations' average price and average sales volume is consistent with consumers preferring stations with a reputation for lower prices.

Our traffic data also reveal that drivers pass enough stations to allow them to search without deviating from their travel routes, as we assume later in the structural model. All stations in our sample display their prices on large signs, so passing traffic in both directions can easily observe prices. Panel (b) and (c) of Figure 3 show the distributions of the number of prices drivers see as well as the number of stations they directly drive past

similar to the correlation of 0.3 reported by Houde (2012).

²³An island is an elevated platform where pumps are located. The number of islands provides a better measure on the station scale than the number of pumps. A small station can have four or more pumps cramped on an island, whereas a large-format station generally has two pumps sitting on an island.

along their travel routes. On average, a driver sees 3.5 prices and directly drives past 2.2 stations along their travel route. Thus, the number of options for drivers is comparable to the number of stores searched by consumers before purchases in other markets documented by the search literature.²⁴

Two Types of Price Uncertainty

Frequent price changes in the retail gasoline market make it difficult for consumers to maintain accurate price information. Figure 4 shows the distribution across days of the proportion of stations changing their price from the previous day. The average probability of such a price change is 32 percent.²⁵ Frequent price changes generate two types of uncertainty in the market: (i) ex-ante uncertainty about the price at each station and (ii) uncertainty about the overall price level in the market. To further analyze the different sources contributing to these two types of uncertainty, we perform the following regression,

$$p_{jt} = \sum_{j=2}^J \psi_j \text{Station}_j + \sum_{t=1}^T \gamma_t \text{Day}_t + \nu_{jt}, \quad (1)$$

where the price at station j on day t is decomposed into a station fixed effect ψ_j , a day-of-sample fixed effect γ_t , and an idiosyncratic error term ν_{jt} .

Using this decomposition, overall variation in price can be viewed as the sum of persistent price differences across stations, captured by station fixed effects ψ_j and a time-variant component $\tilde{p}_{jt} = \gamma_t + \nu_{jt}$. The time-variant price, \tilde{p}_{jt} , combines the day-of-sample fixed effect, γ_t , which is driven by changes in aggregate market conditions (for example, wholesale cost) common to all stations, and the station-day-specific price shock, ν_{jt} .

Drivers regularly observe prices during everyday driving and are likely to be relatively knowledgeable about which stations tend to have higher or lower prices. In contrast, drivers are largely unaware of the wholesale market conditions responsible for fluctuations in the average retail price of gasoline (except what they might infer from recently observed prices),

²⁴Using data on individual online browsing and purchase histories, De los Santos et al. (2017) find consumers visit on average 2.82 online retailers before buying an MP3 player, and De los Santos (2018) finds consumers searched 1.3 online bookstores before purchasing a book.

²⁵This number is only a conservative measure of the proportion of stations with price changes on a day due to missing prices. Some price changes are likely not recorded.

and, on any given day, they don't know which stations are charging unusually high or low prices until they begin searching. Our empirical model captures these institutional features by assuming that average price differences across stations, reflected by station fixed effects ψ_j , are known by consumers so that their uncertainty over prices results entirely from temporal price variation. This structure allows for the possibility that customers respond more strongly when a station begins consistently offering a lower price (which becomes known to consumers) than when a station offers an abnormally low price on a given day (which would not be known prior to search).

Fluctuations in station prices relative to one another are represented by the station-day-specific price shocks, ν_{jt} (e.g., Lewis 2008). Chandra and Tappata (2011) provide direct empirical evidence of such variation, showing that gasoline station pairs exhibit reversals from their normal price ordering approximately 15 percent of the time. This variation creates uncertainty by preventing consumers from knowing a station's location in the current price distribution prior to search. Empirical studies that structurally estimate models of consumer search focus primarily on this type of uncertainty. The known price distribution assumption assumes that market conditions are constant over time and known to consumers (e.g., Hong and Shum 2006; Wildenbeest 2011; Nishida and Remer 2018).

In reality, however, frequently changing market conditions can make it difficult for consumers to know the level of the overall price distribution, as measured by γ_t . In our sample, the price level fluctuations are responsible for 93 percent of the overall retail price variation. This introduces a second important source of price uncertainty. In this environment, consumers may form expectations of the price level today based on prices observed during past trips or gasoline purchases, resulting in biased beliefs (Lewis 2011). Collecting new price observations allows consumers to learn more about the current distribution, but the presence of station-specific price variation, ν_{jt} , prevents them from fully resolving uncertainty in γ_t . Therefore, biased prior beliefs will continue to impact consumers' posterior beliefs about the current price level, though the weight placed on these priors decreases as more new information is obtained. Incomplete knowledge of the price distribution can have important impacts on consumer behavior. For example, biased beliefs provide one explanation for why wholesale gasoline cost increases are often passed through to retail

prices more quickly than cost decreases, as depicted in Figure 1.

Based on the decomposition in Equation 1, Table 2 presents the relative magnitudes of the two components of price variation that give rise to uncertainty. The ν_{jt} , has a standard deviation of 3.2 cpg, confirming the presence of substantial idiosyncratic price variation within a day. In addition, the price levels, γ_t , exhibit considerable fluctuation during our sample period, spanning from a minimum of \$1.08 to a maximum of \$2.04 before taxes. The average across days of the absolute difference in the price level, $|\gamma_t - \gamma_{t-1}|$, is 0.8 cpg. However, the discrepancies between the priors and the actual price levels might be much larger, as consumers are likely to use the price at their last gasoline purchase as a reference price (Lewis, 2011). For example, the average absolute difference between the current price and the price 7 days prior is 4.5 cpg.

We now provide some descriptive evidence on whether consumers know about the current price distribution under the premise of consumer search. If consumers have correct knowledge about the actual price distribution, past prices should not affect consumers' search decisions. In contrast, when consumers are uncertain about the actual price distribution and formulate their price expectations based on prices acquired from past driving trips or purchases, these past prices may bias consumers' perceived benefit of price search and influence consumer search in certain directions. In particular, lower past prices may bias consumers' price expectations downward, causing more consumers to postpone their purchases to future trips searching for better prices, consequently lowering current gasoline sales. Similarly, higher past prices may reduce the perceived benefit of search or postponement, leading to more purchases on current driving trips. To investigate the relationship between purchases and past prices, we regress the logarithm of a station's daily transaction volume on its own price, its closest competitors' prices, and the average price level in the city 7 days prior while controlling for station as well as day-of-week and month-of-sample fixed effects.²⁶ We measure a station's closest competitors in terms of the amount of traffic the stations have in common. Note that under imperfect price information, price changes at subsequent stations along a travel route are unknown to consumers and thus do not affect

²⁶We use the price level 7 days ago, γ_{t-7} , as a measure of the prices that consumers observe on their past trips or purchases. We have estimated the model using the price level on various days prior and obtained similar results.

their search or purchase decisions at the current station. In other words, a price change at an upstream station can influence the demand at a downstream station but not the other way around. Therefore, we define a station’s common traffic shared with a neighbor station as the proportion of the given station’s passing traffic (in all directions) that has previously passed the neighbor. We then rank neighbors in terms of common traffic shares for each station.²⁷ Table 3 provides the coefficient estimates for the panel regression. An increase in a station’s daily gasoline sales predicts a decrease in its own price and an increase in the prices of its closest two neighbors, as expected. Importantly, the coefficient on the logarithm of the past price level is positive and precisely estimated in all specifications. These results are consistent with consumers being uncertain about the current price level but contradict the assumption of a known prior distribution, where past prices should not affect search and purchase decisions.

It is also possible that past prices influence current demand not through prior beliefs but due to consumers timing their gas purchases using their gas tank storage. If more (fewer) consumers fill up their tank as a result of recent low (high) prices, fewer (more) will be in the market to purchase gas today. To investigate this potential channel formally, Column (3) introduces the logarithm of past total gasoline transactions from various days ago to control for the number of consumers who have purchased gas in the recent past. If the negative lag-price effect is entirely due to purchase timing, then the lag-price coefficient should become a much less important predictor once lagged total transactions are controlled for. In contrast, controlling for recent purchase volumes in Column (3) has minimal impact on the magnitude and estimated precision of the coefficients on past prices, suggesting that past prices affect purchase decisions through prior belief formation.

Moreover, adjustments to purchase timing should generate a negative relationship between past and current sales, whereas, Column (3) shows that the coefficient on the past sales 1-day ago is positive and precisely estimated. The positive coefficient likely reveals some degree of serial correlation in unexplained shocks to gasoline demand. The coefficient on past sales from 2-days ago is negative and precisely estimated, consistent with timing purchases, but the coefficient estimates become imprecise for longer lags. The estimates

²⁷The magnitude of these common traffic shares is discussed later when we consider station pair characteristics in Section 7.

suggest that adjustments to purchase timing, if they occur at all, tend to happen within a window of several days, consistent with the findings in Levin et al. (2022). Most importantly, the coefficient on the lagged price level remains precisely estimated and unaffected by the inclusion of lagged sales volumes, suggesting that past prices primarily impact demand through the formulation of consumers' prior expectations about the price distribution.

5 Model of Sequential Search

Based on the important institutional details of the retail gasoline market, we specify a model of price search in which consumers are uncertain about the current price distribution and learn about the distribution as they observe prices. The model considers heterogeneous and forward-looking individuals, each traveling along a particular (exogenously determined) route, sequentially encountering a known set of stations, perhaps from home to work or from home to a store. Consumers hold prior beliefs about the distribution of prices in the market, likely based on the prices observed from past driving trips or purchases. As a consumer passes each station, she updates her price beliefs before optimally deciding whether to stop and purchase or continue on to potentially purchase at a subsequent station.²⁸ Consumers also have the option of postponing purchase until a future trip but some will incur a postponement cost (that varies across consumers) to reflect that certain consumers need to purchase gasoline more urgently than others.²⁹ Therefore, the probability that a consumer will purchase from a station depends on the realized utility of purchasing at the observed price and the expected value of continuing to search given her posterior belief of the price distribution.

Although our search model characterizes an individual consumer's purchase decision, our empirical model will be estimated using station-level quantity and price data. A station's potential customers may be traveling along many different routes and encountering different sets of competing stations. The quantity of gasoline sold at a particular station

²⁸Although gasoline consumption does respond somewhat to changes in price, this study focuses instead on how consumers decide where and when to make that gasoline purchase. In our empirical estimation, we use time fixed effects to control for the changes in overall demand level as detailed in Section 6.

²⁹We allow for a mass of consumers with zero postponement cost to account for drivers whom have recently purchased gas and are not considering another purchase.

can then be modeled by aggregating individual predicted purchase decisions over the empirical distribution of consumers across search routes. The following subsections detail the different components of the individual search model. Then, in the next section, we discuss the construction of the empirical model, including aggregation to the station level and the additional assumptions required for estimation.

Utility

Gasoline demand is characterized by consumers sequentially searching the prices of stations along their travel routes. We consider a city containing a set, \mathcal{J} , of J stations indexed $j = 1, 2, \dots, J$. Consumers each demand 10 gallons of gasoline.³⁰ We assume consumers have an indirect utility for gasoline at station j on day t equal to:

$$u_{jt} = X_j\beta - p_{jt},$$

where X_j represents station j 's non-price characteristics, and p_{jt} is the unit cost of gasoline (per gallon price multiplied by 10 gallons). The coefficient on the price is normalized to -1 , so utilities are expressed in dollar value. Because around 30 percent of the stations in our sample do not have market share data, it is not possible to introduce station fixed effects or allow unobserved station attributes in the model. Instead, we use brand and scale group dummies to parameterize the station-specific unobserved attributes, similar to the approach used by Goldberg (1995).³¹

Based on the price decomposition in Equation 1, we can rewrite the indirect utility as

$$\begin{aligned} u_{jt} &= X_j\beta - \psi_j - \gamma_t - \nu_{jt} \\ &= V_j - \tilde{p}_{jt}, \end{aligned} \tag{2}$$

where γ_t represents the daily average price level in the city, ψ_j captures the persistent

³⁰A unit of gasoline purchase of 10 gallons is a scalar chosen for the convenience of interpreting the estimation results.

³¹Given that gasoline is less differentiated than most products, any remaining unobserved quality should be negligible. Additionally, we find little correlation between the estimated utility generated by the non-price characteristics $X_j\hat{\beta}$ and the persistent price difference across stations $\hat{\psi}_j$, suggesting that any remaining unobserved quality is unlikely to be correlated with price.

price difference between stations, and ν_{jt} is the idiosyncratic deviation of station j 's price on day t from its own average and the city average. Therefore, the customer's indirect utility of consuming at station j on day t can be partitioned into two components: the value of station j 's time-invariant characteristics, $V_j = X_j\beta - \psi_j$, and a time-varying price component, $\tilde{p}_{jt} = \gamma_t + \nu_{jt}$.

This partition of the customer's indirect utility function is motivated by the features of the retail gasoline market. Repeated observations and frequent purchases at a number of stations allow consumers to become aware of the station characteristics that are constant over time. These include the station's location, brand, and reputation for being a high- or low-priced station (represented in the model by ψ_j). The V_j component then represents the part of utility known to consumers before search. In contrast, time-variant prices, representing the changes in prices over time and across stations, are unknown to consumers, as discussed in the previous section.

Consumer Learning and Prior Belief

We assume that time-variant prices, $\tilde{p}_{jt} \sim N(\gamma_t, \sigma^2)$, where σ^2 denotes the magnitude of the actual price dispersion.³² We make the independence and normality assumptions to make the model tractable because the conjugate prior of a normal distribution is itself, even though the distribution may be inconsistent with the equilibrium price distribution.³³ Based on past experiences, consumers are likely familiar with the typical level of idiosyncratic price variability in the market. Therefore, we assume that consumers know σ^2 .³⁴

Because retail gasoline prices frequently rise and fall in response to volatile wholesale prices, consumers are uncertain about the average price level, γ_t . We capture this uncertainty by assuming consumers hold some common prior beliefs about the price level.³⁵ In

³²We assume that the price dispersion is constant over time. Although the degree of gasoline price dispersion has been shown to fluctuate over time when price levels change (Lewis and Marvel, 2011; Chandra and Tappata, 2011), our article abstracts from this second-order effect.

³³Empirically, the normal distribution provides a fairly accurate characterization of the distribution of ν_{jt} , though the empirical distribution does have somewhat heavier tails. Solving the supply-side prices in the presence of complex route structures and evolving consumer beliefs is beyond the scope of this article.

³⁴This assumption is also necessary for the identification of the prior weight (Mehta et al., 2003; Ursu et al., 2020). In the analysis, σ is set to 0.32 to match the empirical distribution presented in Table 2.

³⁵The common prior assumption is for analytical tractability. However, as consumers observe different prices along different search routes, their posterior beliefs become different.

particular, we assume consumers perceive possible price levels as random variables,

$$m_{0t} \sim N\left(\mu_{0t}, \frac{\sigma^2}{\alpha_0}\right), \quad (3)$$

where μ_{0t} is the mean (expectation) of the perceived price levels, later referred to as the prior mean. Following the literature, we denote the variance of the prior belief as a ratio of the known price variance σ^2 and α_0 . Here, α_0 , commonly known as the prior weight, is inversely related to the prior uncertainty about the price level. For example, a smaller α_0 suggests a more diffuse prior.

Due to relative price variation, each price observation only provides a noisy signal of the true price level, γ_t . As consumers observe new prices, they update their beliefs about the price level. Let x_n be the realization of the n th time-variant price observation. According to Bayes' rule, the posterior belief about the price level after observing n prices follows a normal distribution

$$m_{nt} \sim N\left(\mu_{nt}, \frac{\sigma^2}{\alpha_0 + n}\right), \quad (4)$$

where

$$\mu_{nt} = h(\mu_{n-1,t}, n, x_n) = \frac{(\alpha_0 + n - 1)\mu_{n-1,t} + x_n}{\alpha_0 + n}. \quad (5)$$

The posterior uncertainty, $\frac{\sigma^2}{\alpha_0 + n}$, falls in the number of price observations. Based on Equation 5, the posterior mean of the perceived price level can also be expressed as a weighted average of the prior mean and the new price observations,

$$\mu_{nt} = \frac{\alpha_0}{\alpha_0 + n}\mu_{0t} + \frac{1}{\alpha_0 + n} \sum_{k=1}^n x_k. \quad (6)$$

The posterior belief, which captures the learning process, depends on two critical components: prior uncertainty and prior mean. The prior weight, α_0 , determines the speed of learning. When α_0 is smaller, meaning that consumers are more uncertain about their prior beliefs, the posterior mean is updated more by each price observation. On the other hand, a larger α_0 suggests a slower update, as the posterior mean depends more on the prior mean. Moreover, when consumers are perfectly certain of their prior beliefs about the price level (α_0 is infinite), the posterior mean always equals the prior mean regardless of

the new price observations, $\mu_{nt} = \mu_{0t}$. In other words, no learning occurs, and consumers believe that any observed price deviation from their prior mean is the result of a station’s specific price change rather than a market-level price change.

The prior mean, μ_{0t} , also plays a vital role in formulating the posterior beliefs. Importantly, the prior mean does not need to equal the actual price level, γ_t . In fact, as discussed in Section 4, prior beliefs are likely biased as consumers formulate their prior beliefs based on the prices observed from previous gasoline purchases or driving trips.

Ordered Search

Drivers typically pass multiple gas stations while driving to their desired destinations. As a result, unlike some other product markets, consumers can sequentially search the prices of multiple stations with zero search cost. In practice, drivers rarely alter their routes or make separate trips to visit additional stations. Hence, we adopt a model that assumes such deviations from the travel route, including recall (driving back to a previously passed station), are too costly. Consumers’ price search is sequential and ordered, as they know ex-ante the predetermined order in which they will pass a specific set of differentiated stations.

We assume that consumers update their beliefs based on the price observations from both sides of the street. However, we introduce a visit cost (or turn cost), $\tau_{rn} \in \{0, \tau, \infty\}$ to account the higher cost of visiting stations across the street. This cost is zero if the traveler is on the same side of the street as the station or if a left-turn is easy to make, but will take on a non-negative value, τ , if a left-turn requires waiting for traffic lights at a major intersection. For travelers who are unable to visit the station due to left-turn restrictions this turn cost parameter becomes infinite.³⁶ With some abuse of notation, let $r(n)$ return the station index j for the n th station along route r . This station’s route-specific ex-ante known utility is then

$$V_{rn} = V_{r(n)} - \tau_{rn}. \tag{7}$$

Consider a consumer i ’s search decision as she drives along a route r on day t . For notational simplicity, the day index t is suppressed until necessary. As the consumer drives

³⁶See Section 4 for additional discussion of how left-turn difficulty is determined for each station.

to each station, she costlessly observes the price. She updates her belief about the prices at the other stations before deciding whether to purchase gasoline at this station or go to the next one. This decision amounts to an optimal stopping problem involving a value function, W_r , with three state variables: the number of prices already observed, n , the price at the current station, x_{rn} , and the posterior mean, μ_{rn} . Upon observing the price at a station n prior to the final station on the route, the consumer trades off the realized utility at the n th station with the value of continuing searching, evaluated based on her current estimates of the price distribution given the price information obtained. Therefore, the value function can be recursively defined as,

$$W_{irn}(\mu_{rn}, x_{rn}) = \max \left\{ V_{rn} - x_{rn}, \int W_{irn+1}(h(\mu_{rn}, n+1, x_{rn+1}), x_{rn+1}) \cdot dF_{rn}(x_{rn+1}) \right\}, \quad (8)$$

where $F_{rn}(x_{rn+1})$ is the posterior predictive distribution of possible unobserved prices at the $n+1$ th station given the n prices already observed along route r . As we show in online Appendix A, it follows a normal distribution with $x_{rn+1} \sim N\left(\mu_{rn}, \sigma^2 + \frac{\sigma^2}{\alpha_0+n}\right)$. The predictive distribution takes into account both the station-specific price variation, σ^2 , conditional on a possible price level as well as the posterior uncertainty over the price levels, $\frac{\sigma^2}{\alpha_0+n}$.

In practice, drivers typically travel on a variety of different routes to and from their various destinations. Some may choose not to purchase on their current trip, instead continuing to search for a better deal on a future trip along a different route. This option becomes increasingly costly when a consumer is close to running out of gas. In our model, if a consumer does not purchase from a station along the current travel route, she pays a postponement cost c_i . This c_i will be higher for those who need to purchase now, and lower for those seeking to buy gasoline but not under pressure to do so immediately.³⁷ Because our data do not track individuals' driving behaviors over time, we assume that consumers face the same set of ordered search routes \mathcal{R} . Let λ_r denote the share of drivers traveling along route r given by the traffic data.

³⁷The postponement cost can be interpreted as the psychological cost of worrying about running out of gasoline or expectation of a future stock-out cost.

Therefore, at the final station $n = N_r$, the value function becomes

$$W_{irN_r}(\mu_{rN_r}, x_{rN_r}) = \max \left\{ V_{rN_r} - x_{rN_r}, \right. \\ \left. - c_i + \sum_{r' \in \mathcal{R}} \lambda_{r'} \cdot \int W_{ir'1}(h(\mu_{rN_r}, 1, x_{r'1}), x_{r'1}) \cdot dF_{rN_r}(x_{r'1}) \right\}. \quad (9)$$

The continuation value of search at the end of a search route is then the sum of the postponement cost and the weighted sum of the expected value function at the start of a new travel route. We assume consumers are myopic in the sense that each day the consumer solves a completely new ordered search problem. When considering postponement, consumers perceive the expectation of the future price level to be the same as the expectation of the price level based on their subjective posterior beliefs at the end of a route.³⁸ However, consumers' uncertainty about the future price level is reset to σ^2/α_0 ($n = 0$), as they have not yet observed any prices on the next travel route. Therefore, $F_{rN_r}(\cdot)$ is a normal distribution with mean μ_{rN_r} and variance $\sigma^2 + \frac{\sigma^2}{\alpha_0}$.

Conditional on consumer taste and learning parameters θ , we denote the continuation value of search at any station $n < N_r$ along route r as

$$Z_{rn}(\mu_{rn}, c_i|\theta) = \int W_{irn+1}(h(\mu_{rn}, n+1, x_{rn+1}), x_{rn+1}) \cdot dF_n(x_{rn+1}) \quad (10)$$

$$= \int \max\{V_{rn+1} - x_{rn+1}, Z_{rn+1}(h(\mu_{rn}, n+1, x_{rn+1}), c_i|\theta)\} \cdot dF_{rn}(x_{rn+1}), \quad (11)$$

where Equation 11 is obtained by combining Equation 8 and Equation 10. At the final station $n = N_r$,

$$Z_0(\mu_{rN_r}, c_i|\theta) = Z_{rN_r}(\mu_{rN_r}, c_i|\theta) = \\ - c_i + \sum_{r' \in \mathcal{R}} \lambda_{r'} \int \{V_{r'1} - x_{r'1}, Z_{r'1}(h(\mu_{rN_r}, 1, x_{r'1}), c_i|\theta)\} \cdot dF_{rN_r}(x_{r'1}). \quad (12)$$

Proposition 1 *The continuation value of search can be written as $Z_{rn}(\mu_{rn}, c_i|\theta) = z_{rn}(c_i|\theta) - \mu_{rn}$ for any $r \in \mathcal{R}$ and $n = 1, 2, \dots, N_r$. Solving the recursive relationship presented in Equation 11 and 12, the continuation value of search can be simplified as follows:*

³⁸This is a realistic assumption, as the average price level series follows a random walk.

When $n < N_r$,

$$z_{rn}(c_i|\theta) = z_{rn+1}(c_i|\theta) + \sigma \sqrt{\frac{\alpha_0 + n}{\alpha_0 + n + 1}} \cdot (\zeta_{rn+1} \cdot \Phi(\zeta_{rn+1}) + \phi(\zeta_{rn+1})), \quad (13)$$

and when $n = N_r$,

$$z_{rN_r}(c_i|\theta) = z_0(c_i|\theta) = -c_i + \sum_{r' \in \mathcal{R}} \lambda_{r'} \left[z_{r'1}(c_i|\theta) + \sigma \sqrt{\frac{\alpha_0}{\alpha_0 + 1}} \cdot (\zeta_{r'1} \cdot \Phi(\zeta_{r'1}) + \phi(\zeta_{r'1})) \right], \quad (14)$$

where $\zeta_{rn+1} = \frac{V_{rn+1} - z_{rn+1}(c_i|\theta)}{\sigma \sqrt{\frac{\alpha_0 + n}{\alpha_0 + n + 1}}}$. $\Phi(\cdot)$ and $\phi(\cdot)$ are the CDF and PDF of the standard normal distribution respectively.

Proposition 1 shows that the continuation value of search Z_{rn} is the sum of posterior mean, μ_{rn} , and a function of the postponement cost, z_{rn} , conditional on the consumer parameters (see online Appendix B for proof). In other words, $z_{rn}(c_i|\theta)$ summarizes the value of the time-invariant characteristics of the remaining options along a route for the consumer with postponement cost c_i . Based on the recursive relationship, we can numerically solve for z_0 , and subsequently z_{rn} as a function of c_i at any station along any route. In practice, the solution is given by linear interpolation.

Having not purchased at any previous stations on the route, purchase occurs at the n th station if $V_{rn} - x_{rn} \geq z_{rn}(c_i|\theta) - \mu_{rn}$, where $\mu_{rn} = \frac{\alpha_0}{\alpha_0 + n} \mu_0 + \frac{1}{\alpha_0 + n} \sum_{k=1}^n x_k$. It is straightforward to show that z_{rn} is decreasing in c_i for any $r \in \mathcal{R}$ and $n = 1, 2, \dots, N_r$. Additionally, $z_{rn}(c_i|\theta) > \max\{z_{rn+1}(c_i|\theta), V_{n+1}\}$ for any $n < N_r$ and c_i from Equation 13.

The ordered search generates intuitive properties. First, the value of a route, z_{rn} , increases with the number of stations remaining along the route. Second, z_{rn} is bounded from below by the maximum of the ex-ante known utility of the remaining stations along a route for any c_i . As such, if there is a low-price/high-quality station down the route, where the persistent utility difference between the station and the other stations outweigh the magnitude of any relative price changes, consumers driving along this route will drive past the other stations and buy from the ex-ante desired station.

Therefore, as long as the realized utility at a station is greater than the lower bound of the continuation value of search conditional on the posterior mean, $z_{rn}(+\infty|\theta) - \mu_{rn}$, the critical postponement cost c_{rn}^* that makes the consumer indifferent between purchasing

and continuing to searching satisfies the following,

$$\begin{aligned} z_{rn}(c_{rn}^*|\theta) &= V_{rn} - x_{rn} + \mu_{rn} \\ &= V_{rn} - \frac{\alpha_0 + n - 1}{\alpha_0 + n}x_{rn} + \frac{1}{\alpha_0 + n} \sum_{k=1}^{n-1} x_{rk} + \frac{\alpha_0}{\alpha_0 + n}\mu_0. \end{aligned} \quad (15)$$

If the realized net utility is less than the lower bound, the consumer will continue searching. In this case, c_{rn}^* becomes infinite so that no consumers purchase at the current station.

Intuitively, for the consumer to optimally purchase at a station, her postponement cost must be large enough to make the realized utility greater than the continuation value of search. Therefore, the lower bound of postponement cost necessary for the consumer to stop searching is c_{rn}^* . Additionally, suppose the consumer has already driven past at least one station along the route. For the consumer to optimally purchase at the current station, her postponement cost should not be so large that she has already purchased at a previous station. Therefore, we denote the upper bound of postponement cost as $c_{rn}^{**} = \min(c_{r1}^*, \dots, c_{rn-1}^*)$ when $n > 1$. At the first station, we define $c_{r1}^{**} = \infty$, so that the consumer will purchase if $c_i \geq c_{r1}^*$.

Conditional on searching along route r , the proportion of consumers who purchase from the n th station is

$$q_{rn} = \begin{cases} G(c_{rn}^{**}) - G(c_{rn}^*) & \text{if } c_{rn}^{**} \geq c_n^* \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

where $G(\cdot)$ is the CDF of the postponement costs.

Equation 16 shows that the conditional purchase probabilities along a route are given by the postponement cost distribution evaluated at the critical values. Equation 15 establishes the relationship of the critical values with seller utilities and the posterior beliefs resulting from consumer learning.

6 Estimation

Given our data, two important assumptions are necessary to estimate the structural model. First, we assume that the gasoline transactions are made by a new group of drivers each

day. Therefore, the model is estimated at the day level. The prior mean is parameterized as a weighted average of the price level 7 days ago and the current price level:

$$\mu_{0t} = \pi\gamma_{t-7} + (1 - \pi)\gamma_t. \quad (17)$$

The price level 7 days ago proxies the price observations from recent driving trips or gasoline purchases. The prior bias π , a variable from 0 to 1, reflects the influence of past price observations when consumers formulate their prior price expectations.

This specification of the learning process closely relates to the existing search literature. Due to data limitation, most existing empirical studies on search with learning, including Koulayev (2013) and De los Santos et al. (2017), do not estimate the parameters governing the learning process and assume a correct prior belief ($\pi = 0$) and a prior weight equal to the number of product and seller combinations. Similar to Hu et al. (2019), we estimate the learning process. However, our focus is on how much the past prices bias the prior mean rather than estimating the prior mean itself. Our specifications also allow us to empirically distinguish alternative search model assumptions based on our data. The prior bias π allows us to test empirically whether consumers have correct expectations about the price distribution ($\pi = 0$) or use past price observations as a reference price ($\pi = 1$, e.g. Lewis 2011). When $\pi = 0$ and $\alpha_0 = \infty$, our model nests the standard search models with a known price distribution.

We further aggregate the purchase decisions made by drivers searching along their respective search routes to construct each station’s daily market share, matching the observation level of our gasoline data. We start by defining the total size of the market as the total number of drivers driving on a day who use the financial company’s credit or debit cards for their gasoline purchases. The traffic data describe the average daily drivers in the city. Because other payment methods, such as cash and other companies’ credit or debit cards, are available, we assume a 30 percent market share for the company in the payment means.³⁹ Then, a station’s observed market share is the share of the company’s card users who purchase gasoline at that station on a day. The number of people who purchase at a

³⁹We do not use the actual market share to protect the financial service company’s identity. Different proportionality assumptions do not affect our estimation results.

station is calculated by the daily quantity of gasoline transacted at the station divided by 10 gallons, the unit amount of gasoline per purchase.

Only a small share of the drivers purchase gasoline on any given day because most already have sufficient gasoline remaining in their tanks. In particular, 8.3 percent of the drivers on an average day purchase at the stations where quantity data are observed. To capture the behaviors of the drivers who do not consider a purchase in our model, we allow the distribution of postponement costs to have a probability mass of $1 - \eta_t$ at zero.⁴⁰ The remaining η_t share of drivers have a positive probability of purchasing gasoline on their search routes. Note that η_t includes a set of day of week and month of sample dummy variables to account for the changes in overall demand (frequency of purchase) for gasoline over time.⁴¹ For example, larger η_t reflects more frequent gas purchases during summer seasons. We assume the positive postponement costs follow a log-normal distribution with $E(\ln(c)) = \mu_c$ and $\text{Var}(\ln(c)) = \sigma_c^2$. Therefore, following Equation 16, the share of total consumers who travel on route r and purchase from the n th station in period t can be rewritten as

$$q_{rnt} = \begin{cases} \eta_t \left[\Phi\left(\frac{\ln(c_{rnt}^{**}) - \mu_c}{\sigma_c}\right) - \Phi\left(\frac{\ln(c_{rnt}^*) - \mu_c}{\sigma_c}\right) \right] & \text{if } c_{rnt}^{**} > c_{rnt}^* \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

Let $r^{-1}(j)$ return the location index n of station j on route r . We obtain the expected market share of station j at time t given all station characteristics at time t by aggregating the conditional purchase probabilities across all the routes the station is on,

$$E(s_{jt} | \tilde{p}, X, \lambda) = \sum_{r \in \mathcal{R}} \lambda_r q_{r(r^{-1}(j))t} \mathbf{1}(j \in r). \quad (19)$$

where \tilde{p} is the vector of time-variant prices for all stations, X is a vector of all station characteristics, including the persistent price differences across stations ψ , and λ is the vector of the share of drivers traveling on each possible search route.

⁴⁰Because the price distribution follows a normal distribution and is not bounded, the value of continuing to search is infinite when the postponement cost is zero. Consequently, the drivers with zero postponement cost will never make a purchase.

⁴¹The month-of-sample fixed effects can also account for changes in the market size over time, such as the possible growing adoption of card payments.

Equation 19 shows how we map the observed market share at a station into the conditional purchase probabilities resulting from consumers' search and learning decisions. This relationship allows us to estimate the model using nonlinear least squares. More specifically, to estimate the model, we first use the regression results from Equation 1 to separate prices into persistent price differences across stations ψ_j and time-variant prices \tilde{p}_{jt} . Whereas consumers know ex-ante a station's price reputation along with other station characteristics, they search to realize the time-variant price at each station and update the posterior mean. Second, we calculate the purchase probability at the station-route level. Given a set of parameters, we can numerically solve for $z_0()$, an unknown function of c_i , over a set of equations described by Equation 14 and 13 using the fixed point algorithm. Then we can obtain a numerical solution for $z_{rn}()$ as a function of c_i for each station along each route. All solutions are given by linear interpolation on a grid of c_i . For each station along a route, we can calculate the realized utility net of the posterior mean and determine the critical c_{rn}^* that makes an individual indifferent between purchase and searching using interpolation on the grid of $z_{rn}()$. Equation 16 specifies the probability that consumers on a route will purchase at a station given these critical postponement cost levels. Finally, we aggregate the conditional purchase probabilities over the empirical distribution of search routes to obtain the expected market shares according to Equation 19. We choose the set of parameter values to minimize the sum of squared deviations between the observed and expected market shares.

Identification

The parameters to be estimated are consumer preferences, β s, the postponement cost, μ_c and σ_c , the prior belief, π and α_0 , and a set of time fixed effects in η_t . The joint variation between market shares and station characteristics identifies the preference parameters. For example, given our knowledge of the sampling probabilities observed from the traffic data, a higher-priced station may have a similar market share to a lower-priced station, suggesting that its characteristics are favorable enough to consumers to offset the persistent price difference. The panel feature of the dataset identifies the postponement cost parameter. If prior weight, α_0 , is known, the day-to-day price fluctuations change the critical value

c_{jt}^* at a station. The associated variation in the station's market share can inform the postponement cost density at the critical values and thus μ_c and σ_c .

The identification of the prior weight, α_0 , relies on the exogenous search orders provided by the traffic data. The substitution patterns predicted by search models with learning differ from those where no learning occurs. In particular, negative cross-price elasticities can arise between spatially dispersed stations as a result of the information conveyed through ordered price observations.

To illustrate, consider the following route that consists of three stations labeled 1, 2, and 3, as depicted in Figure 5. For consumers traveling east, the critical postponement cost solves $c_1^* = z_1^{-1}(V_1 - \frac{\alpha_0}{\alpha_0+1}x_1 + \frac{\alpha_0}{\alpha_0+1}\mu_0)$, $c_2^* = z_2^{-1}(V_2 - \frac{\alpha_0+1}{\alpha_0+2}x_2 + \frac{1}{\alpha_0+2}x_1 + \frac{\alpha_0}{\alpha_0+2}\mu_0)$, and $c_3^* = z_3^{-1}(V_3 - \frac{\alpha_0+2}{\alpha_0+3}x_3 + \frac{1}{\alpha_0+3}(x_1 + x_2) + \frac{\alpha_0}{\alpha_0+3}\mu_0)$, respectively. Consider an example where the critical postponement cost at each station follows $+\infty > c_1^* > c_2^* > c_3^*$, so that a positive share of consumers will purchase from each station, and the respective shares are $q_1 = 1 - G(c_1^*)$, $q_2 = G(c_1^*) - G(c_2^*)$, and $q_3 = G(c_2^*) - G(c_3^*)$. If price increases at station 1, c_1^* increases as a result, and the marginal consumers would substitute to station 2. Moreover, through learning, both c_2^* and c_3^* decrease as consumers adjust their (posterior) beliefs upward, expecting a general price increase in the market. In other words, additional search at any subsequent stations becomes less attractive, and consumers will be more likely to purchase at an earlier station along the route. In addition to the marginal consumers gained from station 1, station 2 also gains consumers who would have formerly continued searching and purchased from station 3. Thus, an increase in station 1's price always increases the demand of its closest competitor, which is station 2 in this case. However, the price effect on a subsequent station, like station 3, can be positive or negative. If station 3 loses more consumers to station 2 than those gained from postponement, it has a negative cross-price elasticity of demand with respect to the price at station 1. Where two stations lie within the route structure in our ordered search environment dictates both how closely they compete and how information is conveyed, therefore influencing how a change in one station's price positively or negatively impacts the other's demand. In contrast, standard search models with no learning can only generate non-negative cross-price elasticities. Online Appendix C shows data patterns consistent with the implication of

consumer learning. The magnitude of these negative cross-price elasticities helps determine the size of the learning parameter. Furthermore, observing stations on opposite sides of a street also contributes to the identification of learning. Suppose station 1 were on the other side of the street where drivers passing stations 2 and 3 could not access. If learning is occurring, station 1’s price can affect consumers’ purchase decisions at stations 2 and 3 through posterior beliefs.

Once the prior weight has been identified, we can recover the prior bias π by exploiting the observed correlation between current sales and past price levels. This relies on the assumption that the price level 7 days ago only influences current sales through consumers’ beliefs. As discussed in Section 4, consumers may respond to price fluctuations by adjusting their purchase timing using their gas tank as storage, which can also cause current sales to be correlated with past price levels. However, based on the evidence presented in Section 4 and the findings of Levin et al. (2022), consumers only shift purchases over a shorter period (of 2 to 3 days), and the relationship between current sales and the price level 7-day ago is independent of any changes in sales over the last few days. As such, the data broadly supports our identifying assumption regarding consumers’ beliefs.⁴²

Whether due to purchase timing or persistence in unexplained demand shocks, we do observe some degree of serial correlation in sales that could generate incorrect standard errors if not accounted for in estimation. We address this by using a block bootstrap procedure to calculate standard errors in all our specifications to control for the potential auto-correlation and heteroskedasticity (MacKinnon, 2006). Each block contains 7 consecutive days of observations, allowing dependency in the model’s unobservables across stations and time within a block and randomness of each bootstrap sample.

Notice that we do not include idiosyncratic taste shocks in our model. Similar to Hortaçsu and Syverson (2004), our model relies on heterogeneous postponement costs to provide horizontal differentiation and create non-degenerate market shares. Because gasoline is less differentiated than most products and features frequent price changes, we focus on modeling imperfect price information rather than taste heterogeneity. Moreover, iden-

⁴²In online Appendix D, we also consider a specification where the share of consumers with a positive probability of gas purchase, η_t , depends on lagged sales. Table D.1 suggests that purchase timing and correlation of sales do not confound our estimation results.

tifying postponement costs with taste shocks would require observing data on individual search sequences and variables that shift postponement costs independently of taste shocks, like in Yavorsky et al. (2021). By abstracting away from the possibility of taste shocks, our model demonstrates how exogenous search order can be used to identify the learning parameters using only aggregate data on prices, market shares, and the distribution of search orderings.

Lastly, the fluctuations in the drivers’ overall demand for gasoline over time identify the day of week and month of sample fixed effects in η_t , holding the postponement cost distribution constant. For example, transaction data show that drivers consume more gasoline in the summer than in the winter. Larger η_t estimates in the summer months reflect a larger share of consumers in the market actively searching for gasoline, likely resulting from more frequent purchases.

We present Monte Carlo simulations in online Appendix E to confirm that our estimation approach can separately identify and consistently estimate our model parameters.

7 Results

To facilitate comparison, we estimate our full search model with learning as well as a restricted version that does not incorporate consumer learning (i.e., $\alpha_0 = +\infty$).⁴³ The results are presented in Table 4, with estimates from the full model in Column (1) and estimates from the restricted model in Column (2). Estimates of the bias and learning parameters in Column (1) reveal that consumers’ initial beliefs regarding the distribution of prices are significantly different from the actual price distribution. As consumers observe new prices, they update their beliefs relatively quickly. More specifically, the bias parameter suggests that, before observing any prices, 59 percent of a consumer’s prior belief depends on prices observed in the prior period. This statistically significant weight on the past prices rejects the common assumption that consumers behave as if they have a correct expectation about the price distribution. The average absolute difference between the estimated prior mean and the actual price level is 2.7 cpg, approximately 3.3 times the size of the average day-to-day price change. However, this large initial bias is quickly

⁴³In estimation, we set α_0 to a very large number for the no-learning model.

moderated as the consumer observes new prices along her travel route. The estimated weight on the initial prior is 0.3, suggesting a fast rate of learning. For example, after one new price observation, the bias reduces by 77 percent to 0.6 cpg, and after two new price observations, the bias reduces by 87 percent to 0.4 cpg in expectation. This rapid learning is consistent with the fact that the distribution of gasoline prices changes regularly in response to wholesale cost volatility.⁴⁴ More specifically, the estimated standard deviation of the prior mean is 5.8 cpg, and the standard deviation of the current price and the price 7 days prior is 5.9 cpg. The similarity in these standard deviations suggests that, on average, consumers make their purchase decisions consistent with those who have a decent understanding of the price level variation through their repeated interaction with the market.

With regard to station attributes, the estimates suggest that consumers purchasing 10 gallons of gasoline are willing to pay \$1.11 more to avoid waiting for the left-turn signal at a busy intersection. Consumers also appear to value the features offered at large-format stations, placing a \$0.66 premium on purchasing gas at these stations. In contrast, the willingness to pay for a gas purchase at a small station is \$1.17 lower than at a medium-sized station. We also find that consumers are willing to pay \$0.51 more at a major branded station than at a generic station, all else constant.

Based on the postponement cost estimates, for drivers who have a positive probability of purchasing gasoline, the median cost of postponing purchase to a future trip is \$0.56, and 25 percent of these drivers are willing to pay \$1.15 more to purchase on the current route rather than to postpone their purchase. Note that the postponement cost reflects both the risk of running out of gas as well as the trade-off between purchase on the current trip and purchase on a future trip where a consumer is not only uncertain about the prices and their distribution but also the set of stations she will drive past.

In the restricted version of the search model that does not allow learning (Column (2) of Table 4), the prior weight in Equation 3 is set to positive infinity, meaning that a driver's posterior beliefs about the price level always equal to her prior belief. However, we

⁴⁴The estimated speed of learning is much faster than what is previously assumed by the literature. For example, Koulayev (2013) and De los Santos et al. (2017) choose the prior weight to be the number of product-retailer combinations, resulting in a much slower rate of learning.

still allow for the possibility that past prices may bias the initial prior. Without learning, any bias present will persist and influence purchase decisions regardless of how many new prices a consumer observes. Compared to the full model in Column (1), the search model without learning fits the data worse, as suggested by the pseudo- R^2 .⁴⁵ In addition, the *bias* parameter estimate becomes much smaller and statistically indistinguishable from zero. A comparison to the results of the full model is particularly informative here. When learning is incorporated into the search model, estimates reveal that substantial bias in consumers' priors can arise but is quickly mitigated through learning. In other words, bias may influence a consumer's expectations when visiting the initial stations along the travel route, but will have little impact when visiting subsequent stations. The restricted no-learning model assumes that expectations remain fixed throughout, therefore, making it impossible to identify the presence of biased priors for a subset of stations on the travel route.

The estimates of the postponement costs are also very different for the no learning model. The median postponement cost is \$0.75 from the no-learning model, approximately 33 percent higher than the estimate in the learning model. The higher postponement costs in the no-learning model make sense within our theoretical context (Equation 15). With no learning, consumers behave as if they are certain about the price distribution and how it compares to the current price observation. So this model will predict greater responsiveness to price changes. As a result, the estimated postponement cost parameter will be inflated to allow the no-learning model to fit the relatively low level of price responsiveness observed in the data.

Own-Price Elasticities

We next investigate consumers' predicted responses to station-specific price changes based on the search with learning model as well as the no-learning model. Price elasticities are obtained by simulating how station-specific gasoline purchases change following a one-

⁴⁵The pseudo- R^2 is calculated using $1 - \frac{\sum (s - \hat{s})^2}{\sum (s - \bar{s})^2}$, where \hat{s} is the model predicted market share. Additionally, the $F_{1,15944}$ statistic for restricting α_0 to a very large number is 1067.2, which implies that the null hypothesis of no-learning model is overwhelmingly rejected.

cent increase in a particular station’s price.⁴⁶ Each station’s price in our model can be decomposed into a time-varying component \tilde{p}_{jt} and a price reputation component ψ_j which remains fixed for all periods. Therefore, separate elasticities of demand can be constructed for changes in each price component.

Table 5 summarizes each station’s own-price elasticities based on the parameter estimates in Table 4 for the learning model (Column (1)) and the no-learning model (Column (2)). In the learning model, the average of a station’s own-price elasticity with respect to a change in the time-variant price is -8.4. In contrast, the own-price elasticity with respect to the station’s price reputation is -24.4. In other words, consumers are approximately 2.9 times more responsive to a change in price reputation than to a change in time-variant price. Two factors contribute to the considerable difference in price elasticities. First, a change in the time-variant price is unknown to consumers prior to search, whereas a change in the price reputation is known ex-ante. Consequently, an increase in the time-variant price at a station can only affect the purchase decisions for consumers who have driven by the station and have not purchased from a previous station. On the other hand, an increase in the station’s price reputation may cause more consumers to purchase at earlier stations along their travel route, even before passing that station. Second, when consumers are uncertain about the current price level in the market, a relative price change at a station is confounded by changes in price levels, reducing consumers’ responsiveness. They will be less likely to substitute away from a station charging an unexpectedly high price because of the possibility that it reflects an increase in the entire price distribution rather than a relative increase in the station’s price. In contrast, consumers will respond more strongly to an increase in a station’s price reputation, knowing that it represents a relative deviation from the broader price distribution.

The importance of learning is also highlighted by comparing with the own-price elasticities from the no-learning model in Panel (b). In this model, the average demand elasticity with respect to price reputation is only 22 percent larger than the elasticity with respect to time-variant prices. When consumers do not learn, they do not adjust their prior beliefs

⁴⁶A large proportion of the highway stations’ gasoline transactions likely comes from the outside drivers driving past the city via the interstate highway. Because we do not model these passing drivers’ purchase decisions, we exclude the highway stations from the own- and cross- elasticity analysis.

about the price level as they observe new prices. Therefore, any observed price change at a station is believed to be specific to that station.

Spatial Competition and Cross-Price Elasticities

Estimating a structural model of search with learning also provides a useful framework for examining the nature of spatial competition in the market. The stations in our sample exhibit substantial variation in both their characteristics and their locations within the route network. These differences generate considerable variation in own-price elasticity across stations. The estimated station-average own-price elasticities reported in Table 5 Panel (a) range from -24.40 to -2.67, with a standard deviation of 4.68.⁴⁷ Stations with very elastic demand tend to face competition from similar stations located nearby. In contrast, stations with the least elastic demand often share little common traffic with other stations or have very different characteristics.

Table 6 provides a more complete picture of the degree of spatial differentiation between each pair of stations in our sample. In addition to the estimated cross-elasticity between station pairs, $\widehat{\frac{\partial Q_i}{\partial p_j} \frac{p_j}{Q_i}}$, summary statistics are also reported for the driving distance and the share of common traffic between the stations. Recall from Section 4 that we define Common Traffic as the proportion of station i 's passing traffic (in all directions) that has previously passed station j on their travel routes. Drivers driving along travel routes where station j is downstream to station i are not included in this Common Traffic calculation because price changes at station j are unknown to them when visiting station j .

Not surprisingly, most station pairs have virtually zero cross-price elasticities. After all, more than half of all station pairs have less than 1 percent common traffic share. Only 10 percent of the station pairs, typically involving a station's 4 or 5 closest competitors, have cross-price elasticities larger than 0.16. This makes sense given that 90 percent of station pairs are over two miles away from each other and have a common traffic share of less than 16.0 percent. However, some stations do compete intensively. The top 2.5 percent of the pairs have cross-price elasticities above 1.53 and have common traffic shares of greater than

⁴⁷Wang (2009) finds similar station level price elasticity. He estimates an own-price elasticity of -18.77 for a station located right next to its closest competitor and -6.20 for a station whose closest competitor is 4.2 km away.

35.8 percent.

Importantly, our model suggests that consumer learning can generate negative cross-price elasticities across competing stations, as described in Section 6. Indeed, 17.5 percent of all station pairs are estimated to have negative cross-price elasticities in our learning model. This occurs when consumers adjust their price beliefs upon observing a price change near the start of a route and change their purchase decisions at subsequent stations.⁴⁸ Additionally, for a subsequent station C to have a negative cross-elasticity of demand with respect to an upstream station A , there has to be at least one station in between them. Consumers may interpret a price increase at the upstream station A as a sign that all prices are high and search less. Consequently, consumers who would have formerly purchased at station C may now purchase from an earlier station B prior to reaching station C . On the other hand, if a low price is observed at station A , signaling a potential overall market-wide price decrease, consumers who would have formerly purchased at the in-between station B may now keep searching and purchase at station C . Consistent with the theory, the complimentary station pairs identified by our model are never direct neighbors and are, on average, 3.7 stations away from each other.⁴⁹ In contrast, the no-learning model generates all non-negative cross-price elasticity estimates.

Next, we examine how the estimated cross-elasticities between stations vary with geographic and product differentiation. Whereas most studies of gasoline competition rely on simple measures like driving distance to account for geographic differentiation, our traffic flow data allow us to more directly capture connectedness within the travel network using the amount of traffic the stations have in common. The similarity in station characteristics is also likely to influence substitution patterns. In our search model, the ex-ante known mean utility of a station captures its expected attractiveness, reflecting both its characteristics and average price level. If the mean utilities of two stations are sufficiently different, price changes are unlikely to change the stations' utility ranking on a particular day. As a result, consumers driving past these two stations are unlikely to purchase from the less desirable station even when its price is unexpectedly low.

In Table 7, estimated cross-price elasticities for each station pair are regressed on the

⁴⁸Online Appendix C provides a detailed discussion of the theory of negative cross-price elasticity.

⁴⁹The average is weighted by the number of consumers driving past each station pair.

absolute difference in mean utility of the stations, $|\hat{V}_i - \hat{V}_j|$, as well as various measures of the stations' proximity within the travel network. In all specifications, the absolute difference in mean utility has a precisely estimated negative coefficient, confirming that consumers are more likely to substitute between similar stations. The estimates in Column (1) also show that cross-price elasticities generally decline as the driving distance between stations increases, but this relationship becomes much less precisely estimated once common traffic measures are included (in Columns (2) through (3)). In fact, Column (2) suggests that the traffic share explains a considerable fraction of the variation in substitution patterns between stations. For example, when an additional 10 percent of station i 's passing traffic has previously driven past station j , the cross-price elasticity between the two stations increases by 0.43.⁵⁰

Because some left-turns are costly, the ease with which shared traffic can access station i may impact its cross-price elasticity of demand with respect to an upstream station j . For this reason, we decompose our Common Traffic measure into two variables based on the traffic's ease of access to station i . Common Traffic Easy Access measures the share of station i 's passing traffic that has previously passed station j and can visit station i with no additional cost; that is, station i is on the same side as the traffic or is on the opposite side with an easy left-turn. Correspondingly, Common Traffic Costly Left-Turn measures the share of station i 's passing traffic that has previously passed station j and can only visit station i by a costly left-turn. Indeed, the regression result in Column (3) suggests that cross-price elasticities between stations are significantly higher when the common traffic does not have to make a costly left-turn to access station i .⁵¹ Given the share of common traffic, station i 's cross-price elasticity with respect to an upstream station j is about three times larger when the common traffic can visit station i easily than with a left-turn cost.

8 Biased Priors and Asymmetric Search

Lewis (2011), Yang and Ye (2008), and Tappata (2009) each present theoretical models il-

⁵⁰Many of these patterns are similar to those reported by Houde (2012) in his traffic-based analysis of spatial competition between gas stations.

⁵¹We also considered the easy of access for the price-change station j . However, we did not find statistically differentiated effects.

lustrating why cost increases may be passed through more quickly than cost decreases when searching consumers do not know the true price distribution. This asymmetric pass-through arises because consumers search more intensively when prices rise and less intensively when prices fall. Lewis (2011) and Lewis and Marvel (2011) offer empirical evidence consistent with these predictions. As search intensity increases, competition becomes more intense, and station-level demand becomes more elastic, explaining why gas station margins tend to be low when prices are rising and high when prices are falling (Lewis and Marvel, 2011). Estimating a structural model of search with learning allows us to more systematically demonstrate the mechanisms through which imperfect knowledge of the price distribution generates the asymmetric responses in search and demand elasticity that have been shown to influence margins and cost pass-through.

Using the estimates from our learning model, we construct two measures of search intensity: the share of searching consumers who choose to buy from their current travel route rather than postpone their purchase and the average number of stations searched by consumers who do purchase.⁵² Then, to illustrate how biased prior beliefs impact search behavior, we regress each measure of search intensity on the difference between the current price level and the price level 7 days ago. Column (1) and (2) of Table 8 report the results.

When the current price is unchanged from the previous week, our estimates suggest that 60.7 percent of searching consumers will choose to purchase somewhere along their travel route. On average, these purchasing consumers will observe the prices of 2.6 stations before purchasing. However, when the current price level is 10 percent (or approximately 20 cpg) lower than the previous week, consumers observe the prices of only 2.2 stations before purchasing—a 14.6 percent decrease in the number of stations searched. Additionally, 66.1 percent of searching consumers choose to purchase somewhere along their route rather than postpone in search of better prices. This reduction in postponement corresponds to a 9.0 percent increase in gasoline demand, which is roughly consistent with the descriptive data patterns described in Section 4 and Table 3. In contrast, when the model is estimated with no learning, it only predicts a 2.6 percent reduction in postponement.

⁵²The term *searching consumers* is used here to refer to those that have a positive postponement cost. Recall that our model allows for a mass of consumers with postponement costs equal to zero as only a portion of consumers are interested in buying gas on a given day.

Within our model, changes in search intensity generate changes in the predicted elasticity of demand faced by stations. In Column (3) of Table 8, we regress simulated own-price elasticities at the station-day level on the difference between the current price and the price level one week prior, controlling for station fixed effects. The estimated coefficient, which we refer to as the *elasticity asymmetry coefficient*, suggests that when the current price level is 10 cpg lower than the previous week, own-price elasticities decrease by 0.86 in absolute value—a 10 percent reduction from the average own-price elasticity of -8.4 .

We further illustrate how fluctuations in own-price elasticity relate to prior bias, π , and the prior uncertainty, α_0 , by using our model to simulate a series of counterfactuals. First, to investigate the importance of biased expectations, we vary the degree to which consumers' priors of the current price distribution are biased toward past price levels. While holding the other parameters constant at their estimated level, the prior bias parameter is assigned various values ranging from $\pi = 0$, where the prior distribution is centered around the actual price level, to $\pi = 1$, where the prior is centered around the previous period's price level. We simulate the predicted own-price elasticities for each prior bias parameter value and then regress these elasticities on the change in the price level from the previous week, mirroring the analysis from Table 8. The elasticity asymmetry coefficient and 95 percent confidence interval from each regression are plotted as a function of the prior bias parameter in Panel (a) of Figure 6. The figure also plots the estimated average own-price elasticity for each prior bias value.

As shown in the figure, the average own-price elasticity stays relatively flat at around -8 as the prior bias parameter varies. In contrast, as prior bias increases, the own-price demand elasticities stations face when prices are rising compared to when prices are falling become more asymmetric, as suggested by the more negative elasticity asymmetry coefficient. As a result, the degree of demand asymmetry as a share of the average own-price elasticity increases in the prior bias. More specifically, when consumers have a rational expectation ($\pi = 0$), there is no significant demand asymmetry.⁵³ On the other hand, when consumers formulate their prior belief entirely on the past prices ($\pi = 1$), the demand asymmetry coefficient is -11.1 . On a day when prices are rising and the current price level is 10 cpg

⁵³The asymmetry is not mathematically zero because the own-price elasticity is simulated based on the observed prices.

higher than last week's price level, the predicted margin would be $-11.1 * 0.1 / -8 = 13.8$ percent lower than the margin on an average day in our sample.

A similar counterfactual analysis can be used to evaluate how the degree of certainty consumers attribute to their prior beliefs asymmetrically impacts own-price elasticity following price increases and decreases. When consumers are more certain about their prior beliefs, they place a greater weight on their priors and less weight on newly observed prices when formulating expectations, leading to greater elasticity of demand. For example, consumers are more likely to purchase when encountering a price they think is low because they will be more certain about its low relative position within the price distribution and place little value in the opportunity to continue learning from additional price observations. In addition, because the prior bias is assumed to remain at its estimated value of 0.59, placing additional weight on one's priors allows this bias to generate more persistent differences in consumers' search behaviors when prices are rising and falling, resulting in a more asymmetric response of own-price elasticities.

In Panel (b) of Figure 6 the average simulated own-price elasticity and the asymmetry in that elasticity are plotted for different values of the prior weight parameter. First, the average own-price elasticity grows in absolute value (implying margins are likely to decrease) when consumers place a higher weight on their prior beliefs. This is consistent with the theory that consumers become more responsive to price changes when they are more certain about their prior beliefs. Additionally, as prior weight increases, own-price demand elasticities appear to become more asymmetric between periods of increasing and decreasing prices, though the relationship is not monotone. More specifically, as prior weight increases from 0.1 to 81, the demand asymmetry changes from -5.7 to -17.1. However, the demand asymmetry as a share of the average own-price elasticity remains relatively stable. As a result, when the current price level is 10 cpg higher than the last week's price level, the estimated own-price elasticities will be 10 percent larger, and the implied margins will be 10 percent lower than on an average day, regardless of the prior uncertainty.

9 Conclusion

This article has estimated a dynamic search model with learning where consumers sequentially search for lower gasoline prices in a predetermined order following their travel routes. We allow consumers to be uncertain about the price distribution and hold prior beliefs that may be biased by prices observed during previous purchases. Traffic flow data are used to construct an empirical distribution of search sequences in the market. This novel approach allows us to identify the consumer learning process, postponement costs, and ex-ante seller differentiation using only market share data. We find that consumers place significant weight on past prices when formulating their prior beliefs. However, consumers are relatively uncertain about these prior beliefs. As a result, any initial bias in consumers' expectations diminishes quickly as they update their price beliefs based on new price observations.

By incorporating the consumer learning process, we relax one of the crucial assumptions of standard search models—the assumption that searching consumers are aware of the true price distribution. Prior uncertainty and prior bias are both essential features in the retail gasoline market, as volatile prices make it difficult for consumers to know the true price distribution with any certainty. Consequently, consumers are likely to formulate their expectations of prices based on prices observed in the recent past. More importantly, we systematically demonstrate how prior uncertainty and prior bias can cause demand elasticities to respond asymmetrically to price increases and decreases. This asymmetric demand response offers an explanation for why firms pass through positive cost changes more quickly than negative cost changes—a widely observed phenomenon that cannot be explained by search frictions alone. Our results suggest that price fluctuations will have a larger and more asymmetric impact on demand elasticities when consumers rely more heavily on past prices in forming their priors and when consumers place a heavier weight on these priors as they search for gasoline along their travel route.

The use of travel patterns to simulate unobserved search sequences is grounded in the observation that consumers are likely to search for and purchase gasoline during everyday driving rather than making dedicated trips to purchase gasoline. In addition to the identification of the consumer learning process, our approach has other advantages. First, it allows

us to introduce ex-ante vertical differentiation of stations without suffering from the curse of dimensionality. Second, we use the observed traffic flows to replace the random sampling assumption, allowing us to estimate more realistic substitution patterns that depend on the amount of traffic stations share. Although the integration of travel patterns in a search model is most relevant to the retail gasoline market, we envision its applications in other markets. In cases where sellers have physical addresses, such as in a shopping mall, travel patterns naturally constrain the search order. Even for sellers without physical addresses, the order of visits can be affected by constraints such as a webpage layout.

Our article opens up several avenues for future research. We think the most important is the modeling of the supply side decision. The pricing equilibrium arising in the ordered search environment is likely to be quite different from the equilibrium of a random-search model. Arbatskaya (2007) develops a price equilibrium for a row of sellers facing consumers who travel in one direction. However, pricing decisions in the retail gasoline market are more complicated, as stations are located on multiple travel paths with consumers driving in different directions and passing different sets of competitors. Consequently, the demand at a station, as the sum of the residual demand along each search route, is kinked. Spatial differentiation, together with imperfect price information, creates interesting price dynamics, which we leave for future work to explore. Also, asymmetric cost pass-through is often regarded as anti-competitive and harmful to consumers. A supply-side model would enable researchers to answer important welfare questions. For example, a counterfactual analysis could examine how much consumers would benefit from being informed about the actual price distribution and, therefore, facing a market with no asymmetric search intensity and no asymmetric cost pass-through.

References

- Arbatskaya, M., “Ordered search.” *The RAND Journal of Economics*, Vol. 38 (2007), pp. 119–126.
- Atkinson, B., “On retail gasoline pricing websites: potential sample selection biases and their implications for empirical research.” *Review of Industrial Organization*, Vol. 33 (2008), pp. 161–175.
- Benabou, R. and Gertner, R., “Search with learning from prices: does increased inflationary uncertainty lead to higher markups?” *The Review of Economic Studies*, Vol. 60 (1993), pp. 69–93.
- Borenstein, S., Cameron, A. C., and Gilbert, R., “Do gasoline prices respond asymmetrically to crude oil price changes?” *The Quarterly Journal of Economics*, Vol. 112 (1997), pp. 305–339.
- Byrne, D. P., “Gasoline Pricing in the Country and the City.” *Review of Industrial Organization*, Vol. 55 (2019), pp. 209–235.
- Byrne, D. P. and de Roos, N., “Consumer search in retail gasoline markets.” *The Journal of Industrial Economics*, Vol. 65 (2017), pp. 183–193.
- Chandra, A. and Tappata, M., “Consumer search and dynamic price dispersion: an application to gasoline markets.” *The RAND Journal of Economics*, Vol. 42 (2011), pp. 681–704.
- Dana, J. D., “Learning in an Equilibrium Search Model.” *International Economic Review*, Vol. 35 (1994), pp. 745–771.
- Davis, P., “Spatial competition in retail markets: movie theaters.” *The RAND Journal of Economics*, Vol. 37 (2006), pp. 964–982.
- De los Santos, B., “Consumer search on the internet.” *International Journal of Industrial Organization*, Vol. 58 (2018), pp. 66–105.

- De los Santos, B., Hortaçsu, A., and Wildenbeest, M. R., “Testing models of consumer search using data on web browsing and purchasing behavior.” *The American Economic Review*, Vol. 102 (2012), pp. 2955–2980.
- De los Santos, B., Hortaçsu, A., and Wildenbeest, M. R., “Search With Learning for Differentiated Products: Evidence from E-Commerce.” *Journal of Business & Economic Statistics*, Vol. 35 (2017), pp. 626–641.
- Ellison, S. F., “Price search and obfuscation: an overview of the theory and empirics.” in *Handbook on the Economics of Retailing and Distribution*, Edward Elgar Publishing (2016).
- Goldberg, P. K., “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry.” *Econometrica*, Vol. 63 (1995), pp. 891–951.
- Hong, H. and Shum, M., “Using price distributions to estimate search costs.” *The RAND Journal of Economics*, Vol. 37 (2006), pp. 257–275.
- Honka, E., “Quantifying search and switching costs in the US auto insurance industry.” *The RAND Journal of Economics*, Vol. 45 (2014), pp. 847–884.
- Honka, E., Hortaçsu, A., and Wildenbeest, M., “Empirical search and consideration sets.” in *Handbook of the Economics of Marketing*, Elsevier (2019).
- Hortaçsu, A. and Syverson, C., “Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds.” *The Quarterly Journal of Economics*, Vol. 119 (2004), pp. 403–456.
- Houde, J.-F., “Spatial differentiation and vertical mergers in retail markets for gasoline.” *The American Economic Review*, Vol. 102 (2012), pp. 2147–2182.
- Hu, M., Dang, C., and Chintagunta, P. K., “Search and Learning at a Daily Deals Website.” *Marketing Science*, Vol. 38 (2019), pp. 609–642.
- Jindal, P. and Aribarg, A., “The importance of price beliefs in consumer search.” *Journal of Marketing Research*, Vol. 58 (2021), pp. 321–342.

- Kim, J. B., Albuquerque, P., and Bronnenberg, B. J., “Online demand under limited consumer search.” *Marketing science*, Vol. 29 (2010), pp. 1001–1023.
- Kim, J. B., Albuquerque, P., and Bronnenberg, B. J., “The probit choice model under sequential search with an application to online retailing.” *Management Science*, Vol. 63 (2017), pp. 3911–3929.
- Koulayev, S., “Search with dirichlet priors: estimation and implications for consumer demand.” *Journal of Business & Economic Statistics*, Vol. 31 (2013), pp. 226–239.
- Koulayev, S., “Search for differentiated products: identification and estimation.” *The RAND Journal of Economics*, Vol. 45 (2014), pp. 553–575.
- Levin, L., Lewis, M. S., and Wolak, F. A., “Reference dependence in the demand for gasoline.” *Journal of Economic Behavior & Organization*, Vol. 197 (2022), pp. 561–578.
- Lewis, M. S., “Price dispersion and competition with differentiated sellers.” *The Journal of Industrial Economics*, Vol. 56 (2008), pp. 654–678.
- Lewis, M. S., “Asymmetric price adjustment and consumer search: An examination of the retail gasoline market.” *Journal of Economics & Management Strategy*, Vol. 20 (2011), pp. 409–449.
- Lewis, M. S. and Marvel, H. P., “When do consumers search?” *The Journal of Industrial Economics*, Vol. 59 (2011), pp. 457–483.
- Lewis, M. S. and Noel, M., “The speed of gasoline price response in markets with and without Edgeworth cycles.” *Review of Economics and Statistics*, Vol. 93 (2011), pp. 672–682.
- Lin, H. and Wildenbeest, M. R., “Nonparametric Estimation of Search Costs for Differentiated Products: Evidence from Medigap.” *Journal of Business & Economic Statistics*, Vol. 38 (2020), pp. 754–770.
- MacKinnon, J. G., “Bootstrap methods in econometrics.” *Economic Record*, Vol. 82 (2006), pp. S2–S18.

- Manuszak, M. D. and Moul, C. C., “How far for a buck? tax differences and the location of retail gasoline activity in southeast chicagoland.” *The Review of Economics and Statistics*, Vol. 91 (2009), pp. 744–765.
- Matsumoto, B. and Spence, F., “Price beliefs and experience: Do consumers beliefs converge to empirical distributions with repeated purchases?” *Journal of Economic Behavior & Organization*, Vol. 126 (2016), pp. 243–254.
- McCall, J. J., “Economics of Information and Job Search.” *The Quarterly Journal of Economics*, Vol. 84 (1970), pp. 113–126.
- Mehta, N., Rajiv, S., and Srinivasan, K., “Price uncertainty and consumer search: A structural model of consideration set formation.” *Marketing Science*, Vol. 22 (2003), pp. 58–84.
- Miller, N. H. and Osborne, M., “Spatial differentiation and price discrimination in the cement industry: evidence from a structural model.” *The RAND Journal of Economics*, Vol. 45 (2014), pp. 221–247.
- Moraga-González, J. L., Sándor, Z., and Wildenbeest, M. R., “Consumer search and prices in the automobile market.” *The Review of Economic Studies*.
- Moraga-González, J. L. and Wildenbeest, M. R., “Maximum likelihood estimation of search costs.” *European Economic Review*, Vol. 52 (2008), pp. 820–848.
- Nishida, M. and Remer, M., “The Determinants and Consequences of Search Cost Heterogeneity: Evidence from Local Gasoline Markets.” *Journal of Marketing Research*, Vol. 55 (2018), pp. 305–320.
- Noel, M. D., “Retail gasoline markets.” in *Handbook on the Economics of Retailing and Distribution*, Edward Elgar Publishing (2016).
- Peltzman, S., “Prices rise faster than they fall.” *Journal of Political Economy*, Vol. 108 (2000), pp. 466–502.

- Remer, M., “An empirical investigation of the determinants of asymmetric pricing.” *International Journal of Industrial Organization*, Vol. 42 (2015), pp. 46–56.
- Rothschild, M., “Searching for the lowest price when the distribution of prices is unknown.” *Journal of Political Economy*, Vol. 82 (1974), pp. 689–711.
- Smith, H., “Supermarket choice and supermarket competition in market equilibrium.” *The Review of Economic Studies*, Vol. 71 (2004), pp. 235–263.
- Stigler, G. J., “The economics of information.” *Journal of Political Economy*, Vol. 69 (1961), pp. 213–225.
- Tappata, M., “Rockets and feathers: Understanding asymmetric pricing.” *The RAND Journal of Economics*, Vol. 40 (2009), pp. 673–687.
- Thomadsen, R., “The effect of ownership structure on prices in geographically differentiated industries.” *RAND Journal of Economics*, pp. 908–929.
- Ursu, R. M., Wang, Q., and Chintagunta, P. K., “Search duration.” *Marketing Science*, Vol. 39 (2020), pp. 849–871.
- Wang, Z., “Station level gasoline demand in an Australian market with regular price cycles.” *Australian Journal of Agricultural and Resource Economics*, Vol. 53 (2009), pp. 467–483.
- Weitzman, M. L., “Optimal Search for the Best Alternative.” *Econometrica*, Vol. 47 (1979), pp. 641–654.
- Wildenbeest, M. R., “An empirical model of search with vertically differentiated products.” *The RAND Journal of Economics*, Vol. 42 (2011), pp. 729–757.
- Yang, H. and Ye, L., “Search with learning: understanding asymmetric price adjustments.” *The RAND Journal of Economics*, Vol. 39 (2008), pp. 547–564.
- Yavorsky, D., Honka, E., and Chen, K., “Consumer search in the US auto industry: The role of dealership visits.” *Quantitative Marketing and Economics*, Vol. 19 (2021), pp. 1–52.

Table 1: Summary Statistics of the Station Characteristics

	Obs.	Mean	SD	25%	50%	75%
<i>Panel (a): All Stations</i>						
Avg. Price (\$)	46	1.61	0.06	1.58	1.59	1.68
Avg. Quantity (gal.)	33	978.09	1207.78	235.76	396.09	1525.02
Major Brands	46	0.37	0.49	0.00	0.00	1.00
Number of Islands	46	3.59	1.73	2.00	3.00	5.00
Easy Left-Turns	46	0.26	0.44	0.00	0.00	0.75
No Left-Turns	46	0.28	0.46	0.00	0.00	1.00
Direct Traffic (1,000s)	46	11.52	4.99	8.32	10.66	15.09
<i>Panel (b): Small Stations</i>						
Avg. Price (\$)	25	1.63	0.06	1.59	1.61	1.69
Avg. Quantity (gal.)	19	379.58	372.86	206.64	253.93	376.10
<i>Panel (c): Large-Format Stations</i>						
Avg. Price (\$)	5	1.57	0.01	1.56	1.57	1.58
Avg. Quantity (gal.)	5	2947.87	1508.80	2217.85	2239.35	2572.08

Table 2: Summary Statistics of Relative Price and Price Level

	Obs.	Mean	SD	Min	Max
<i>Relative Price Changes</i>					
ν_{jt}	23732	0.000	0.032	-0.167	0.169
<i>Price Level Changes</i>					
γ_t	529	1.622	0.249	1.084	2.042
$abs(\Delta\gamma_t)$	528	0.008	0.012	0.000	0.108
$abs(\Delta\gamma_{t-7})$	522	0.045	0.037	0.000	0.186

Table 3: Descriptive Evidence of Price Level Uncertainty

	(1)	(2)	(3)
ln(Own Price)	-2.632 (0.147)	-2.961 (0.160)	-2.942 (0.163)
ln(Past Price Level)	1.220 (0.108)	1.132 (0.109)	1.133 (0.116)
ln(1st Neighbor Price)	1.464 (0.152)	1.171 (0.162)	1.124 (0.163)
ln(2nd Neighbor Price)		0.757 (0.144)	0.786 (0.146)
ln(Total Sales 1-Day Ago)			0.234 (0.043)
ln(Total Sales 2-Day Ago)			-0.076 (0.042)
ln(Total Sales 3-Day Ago)			-0.057 (0.042)
ln(Total Sales 7-Day Ago)			-0.060 (0.038)
R^2	0.928	0.928	0.929
Observations	15985	15985	14960

Notes: The dependent variable is the logarithm of the daily transaction volume at each station. We control for station fixed effects, day of week fixed effects, and month of sample fixed effects in all specifications. Robust standard errors are in parentheses.

Table 4: Estimation Results

	Learning		No Learning	
	Coeff.	SE	Coeff.	SE
<i>Prior</i>				
Bias (π)	0.587	(0.061)	0.027	(0.039)
Learning (α_0)	0.304	(0.116)		
<i>Station Attributes</i>				
Major Brand 1	0.505	(0.041)	-0.065	(0.115)
Retail Brand 1	-0.131	(0.041)	-0.168	(0.043)
Retail Brand 2	-0.064	(0.053)	-0.322	(0.094)
Small-Sized Station	-1.173	(0.043)	-1.719	(0.099)
Large-Format Station	0.662	(0.058)	0.904	(0.052)
Left-Turn Cost	1.106	(0.060)	1.767	(0.152)
<i>Postponement Cost</i>				
Constant (μ_c)	-0.574	(0.063)	-0.288	(0.061)
Standard Deviation (σ_c)	1.055	(0.092)	1.250	(0.095)
Pseudo- R^2	0.888		0.870	

Notes: The number of observation is 15985. The day of week and month of sample fixed effect estimates are omitted from the table. The pseudo- R^2 show the fit of the non-highway stations. Standard errors calculated from 200 block bootstrap samples are in parentheses.

Table 5: Summary Statistics of the Station Average Own-Price Elasticity Estimates

	Obs.	Mean	SD	Min	50%	Max
<i>Panel (a): Learning</i>						
Price ($\tilde{p}_{j,t}$)	37	-8.39	4.68	-24.40	-7.66	-2.67
Price Reputation (ψ_j)	37	-24.40	12.71	-60.74	-20.99	-8.70
<i>Panel (b): No Learning</i>						
Price ($\tilde{p}_{j,t}$)	37	-12.64	5.66	-25.91	-10.92	-5.05
Price Reputation (ψ_j)	37	-15.38	6.56	-29.35	-13.04	-6.83

Table 6: Summary Statistics on Cross-Price Elasticities and Measures of Spatial Differentiation Between Stations

	Obs.	Mean	SD	2.5%	10%	50%	90%	97.5%
Cross-Elasticity	1665	0.148	1.221	-0.229	-0.003	0.000	0.163	1.529
Driving Distance	1665	5.526	2.979	0.883	2.000	5.173	9.647	12.597
Common Traffic	1665	0.050	0.096	0.000	0.000	0.009	0.160	0.358

Table 7: Regression Results of Estimated Cross-Price Elasticities on Distance Measures Between Stations

	(1)	(2)	(3)
Driving Distance	-0.067 (0.015)	0.001 (0.005)	-0.000 (0.006)
Abs. Mean Utility Distance	-0.061 (0.024)	-0.082 (0.021)	-0.093 (0.021)
Common Traffic		4.286 (0.945)	
Common Traffic Easy Access			5.144 (1.167)
Common Traffic Costly Left-Turn			1.748 (0.616)
Constant	0.574 (0.106)	0.008 (0.054)	0.026 (0.057)
R^2	0.03	0.11	0.13
Observations	1665	1665	1665


Notes: The dependent variable is the cross-price elasticity. Robust standard errors clustered at the station level are in parentheses.

Table 8: Search Behavior and Demand and Past Prices

	Share of Purchase (1)	No. Station Searched (2)	Demand Elasticity (3)
$\Delta\phi_{t-7}$	-0.272 (0.008)	1.909 (0.052)	-8.619 (3.971)
Constant	0.607 (0.000)	2.624 (0.003)	
R^2	0.74	0.69	0.23
Observations	522	522	18726

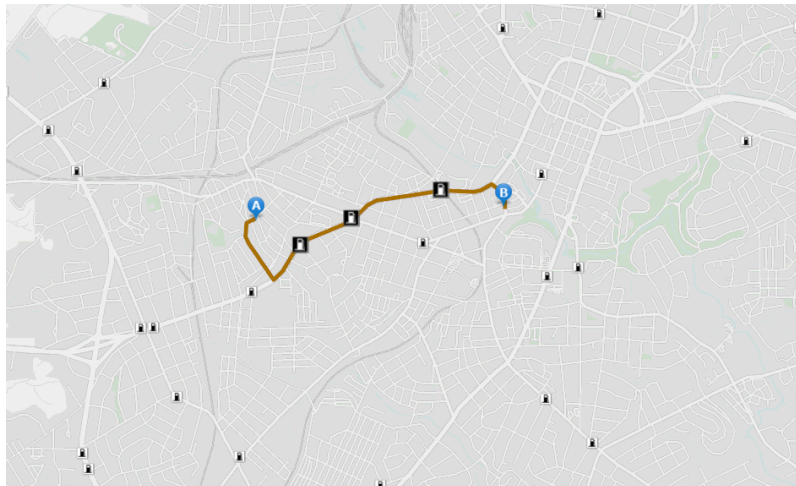
Notes: Robust standard errors are in parentheses.

Figure 1: Average Retail Gasoline Price Level and Wholesale Cost




Note: This figure plots both the average gasoline price in our sample before federal and state taxes are applied and the Gulf Coast regular spot price as a measure of wholesale cost of retail gasoline.

Figure 2: A Travel Route with Stations Passed



Note: This figure represents the road network of a random location. The driving time is less than 5 minutes.

Figure 3: Station Traffic Characteristics



Note: The slope of the log-linear fitted line in Panel (a) is 1.26, significant at 5 percent level.

Figure 4: Proportion of Stations Changing Price From Previous Day




Figure 5: Competing Stations Along a Hypothetical Travel Route


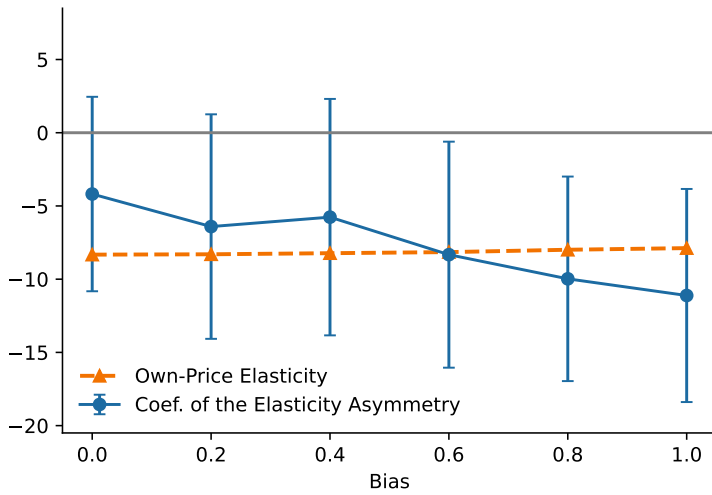


Figure 6: Prior Uncertainty, Prior Bias, and Asymmetry in Own-Price Elasticity

(a) Prior Bias and Elasticity Asymmetry



(b) Prior Uncertainty and Elasticity Asymmetry

