May 2019

Essays on Consumer Search and A+B Auctions

Xiaosong Wu
Clemson University, x.andy.wu@gmail.com

Follow this and additional works at: https://tigerprints.clemson.edu/all_dissertations

Essays on Consumer Search and A+B Auctions

A Dissertation
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
Economics

---

by
Xiaosong (Andy) Wu
May 2019

Accepted by:
Dr. Matthew S. Lewis, Committee Chair
Dr. Babur De los Santos
Dr. Patrick L. Warren
Dr. F. Andrew Hanssen
Dr. Bernardo F. Quiroga

Abstract

This dissertation consists of two chapters. The first paper estimates demand for gasoline in the presence of two types of imperfect price information: ex-ante uncertainty about each station's price and uncertainty about the distribution of all stations' prices. Volatile wholesale cost causes retail gasoline prices to fluctuate regularly, making it difficult for consumers to remain aware of the overall price level in the market or the stations offering the lowest price. In this article, I develop a model in which consumers formulate their prior belief of the current price distribution using the prices observed during past driving trips, and then Bayesian update their beliefs with each new price observed, before deciding whether to purchase gasoline or continue searching for a cheaper price. I estimate this model by utilizing a unique data set of station-level daily gasoline sales and prices, combined with data on the empirical distribution of various traffic flows. My empirical results suggest that consumers are able to learn about the overall price increases or decreases resulting from the wholesale cost movements relatively quickly. In addition, I find that price distribution uncertainty is the primary component of imperfect price information, and if it were eliminated, consumers could achieve 70 percent of the total savings that could be realized by having perfect price information. Furthermore, by incorporating travel patterns, the estimation suggests that cross-price elasticity between two stations depends largely on the amount of common traffic they share.

My second paper studies the effect of complexity in multi-dimensional bidding and competition in A+B (price + quality) auctions using a laboratory experiment. I examine whether the behaviors of human bidders are consistent with the predictions of two alter-

native models of auctions: the Bayes-Nash Equilibrium model and the Quantal Response Equilibrium (QRE) model. I extend the QRE framework to multi-dimensional A+B auctions. The results indicate that the QRE model, as a generalization of the rational models of behavior by allowing decision making errors, predicts bidder behaviors well across different treatments as the number of bidders and the dimensionality of the bid vary.

Dedication

To my parents, Qinghua Tian and Gang Wu, thank you for your unconditional love.

Acknowledgments

I am deeply indebted to my advisor, Dr. Matthew Lewis, for his patience, dedication, insight, and constant guidance and support. I am also grateful for Dr. Babur De los Santos, who constantly encouraged me and challenged me to become a better researcher. I would also like to thank my other committee members, Dr. Patrick Warren and Dr. Andrew Hanssen, for their suggestions and comments regarding my research and presentations. In addition, I would like to thank Dr. Bernardo Quiroga for his continued support and high expectations for me. Thank you for giving me the opportunity to work with you on a joint project that has become the second chapter of my dissertation. I further appreciate the helpful feedback from Dr. Christy Yichen Zhou, Jonathan Ernest, and the participants of the Clemson IO workshop. All errors in this dissertation are mine.

In addition, I would like to thank Dr. William R. Dougan for introducing me to the exciting world of Economics in my first year of Graduate School. I would also like to thank my friends at Clemson for all the laughs and the hours studying together and discussing research. In particular, I thank Yuekai Cheng for being my best roommate.

Finally, I would like to thank my parents, Qinghua Tian and Gang Wu, who love me unconditionally and have dedicated so much of their lives to supporting me and making me the man I am today.

Table of Contents

List of Tables

List of Figures

Chapter 1

Consumer Search with Learning: A Structural Estimation of Gasoline Demand[1]

1.1 Introduction

Many markets are characterized by frequent changes in prices due to rapidly varying market conditions. In these cases, (i) ex-ante uncertainty about each seller's price, and perhaps more importantly, (ii) uncertainty about the distribution of all sellers' prices arise, causing market inefficiencies, including differences in the asking prices among sellers even for homogeneous goods. For example, consumers may not buy from the least expensive seller if it is too costly for them to keep searching for the lowest price or if they are unaware of the existence of this lowest price. When both types of uncertainty are present, consumers engage in costly search not only to realize price offered by sellers (searching), but also to update their beliefs about the current price distribution (learning). While there is rich theoretical and empirical literature that has focused on consumer price search, it has largely ignored

---

[1]This chapter is part of a project that I am working with Frank A. Wolak.

the second type of uncertainty by assuming consumers know the current price distribution.[2] Moreover, the few empirical papers that have incorporated both types of uncertainty have not estimated the consumer learning process.[3]

In this study we investigate consumer search with learning behavior in the context of the retail gasoline market. Two salient features, which are depicted in Figure 1.1, make the retail gasoline market an ideal place for analyzing this phenomenon. First, frequent price changes resulting from a volatile wholesale cost make it difficult for consumers to maintain accurate price information for each station as well as the distribution of these prices.[4] Second, the uncertainty of consumers' beliefs about the current price distribution is often used to explain the asymmetric price response to changes in cost (Lewis, 2011). In particular, when consumers do not know the current price distribution, they may form their initial perception based on prices observed during past driving. In this case, when prices and costs decrease, consumers' beliefs of the price distribution tend to be too high, resulting in low expected gains and less searching, which, in turn, causes higher profit margins for the seller. If consumers follow this behavior, then assuming that they know the current price distribution will lead to upward-biased estimates of search costs (either high search costs or low expected gains from this search can contribute to the low levels of search). Consequently, welfare analysis may overstate the benefits of policies that reduce

---

[2]Pioneer works, such as Stigler (1961) and McCall (1970), among others have established the theoretical foundation for models of consumer search and emphasize the importance of search frictions in explaining imperfectly competitive market behavior under the assumption that consumers know the price distribution. Under this key assumption, a large number of empirical papers have estimated the search costs that contribute to sustained price dispersion in various market, for example, Hong and Shum (2006) on on-line textbooks; Hortaçsu and Syverson (2004) on S&P index funds; Wildenbeest (2011) on supermarkets; Koulayev (2014) on hotel bookings online; Honka (2014) on auto insurance; ? on automobile markets; De los Santos et al. (2012) and De los Santos (2018) on online books; Nishida and Remer (2018) on retail gasoline markets.

[3]Rothschild (1974) was the first to develop a theoretical model in which consumers search and learn about the price distribution in the process. To our knowledge, Koulayev (2013) and De los Santos et al. (2017) are the only two papers that integrate consumer search with learning in empirical work. However, in both papers, parameters that determine the consumer learning process are not estimated but chosen by the researchers.

[4]Although we focus on how changing market conditions cause uncertainty about the price distribution in this paper, inexperience about the market can also contribute to this uncertainty. Using survey data, Matsumoto and Spence (2016) find that college students who have no prior experience in purchasing textbooks online tend to expect online prices to be higher than what are observed empirically, and consumers with more experience in the marketplace generally have more accurate beliefs about the price distribution, which is consistent with learning.

information uncertainty or restrict firm entry in an attempt to lessen the total costs of the search process.

Figure 1.1: Average City Gasoline Price and Cost
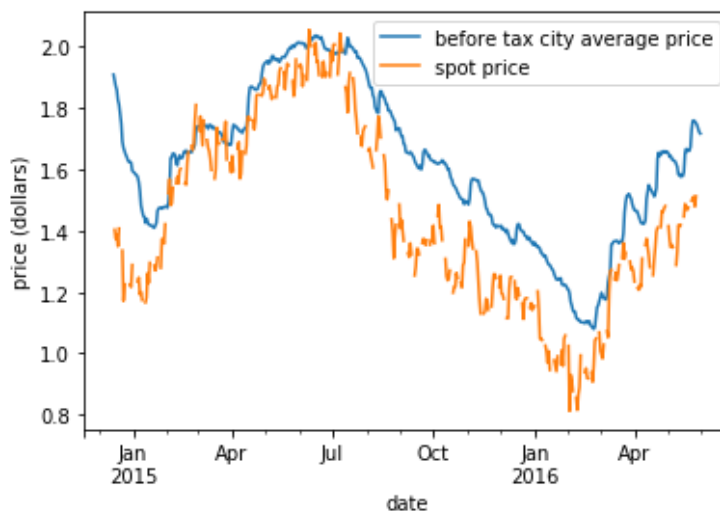
Notes: This figure plots both the average gasoline price in our sample before federal and state taxes are applied and the Gulf Coast regular spot price as a measure of wholesale cost of retail gasoline.

In this paper, we propose a model of search with learning that builds on Rothschild (1974), with an emphasis on spatial and ex-ante vertical differentiation of sellers. We apply the model to a retail gasoline market where consumers are assumed to formulate their prior beliefs on the current price distribution using the prices observed in the past, and then Bayesian update their beliefs with each new price observed along their predetermined travel routes, before deciding whether to purchase gasoline or continue searching. Consumers maximize their utility from a gasoline purchase by comparing if the realized utility at each station after observing its price is higher than the expected value of the remaining stations on the route and the alternative of waiting to purchase during a future trip. If consumers observe a price higher than expected price at a station, they continue driving since they expect to find lower prices at their remaining options on the route. However, learning influences consumers' purchase decisions because as they sample additional high prices, and the newly obtained price information begins to outweigh their prior beliefs, their

posterior beliefs of the actual price distribution are adjusted upward. They are now more likely to make a purchase as they believe that prices at other stations are higher.

To estimate the model, this study uses a novel dataset of daily transactions from each gasoline station in a city from December 14, 2014, to May 31, 2016, and combine it with data on daily prices and an empirical distribution of travel routes. We model a station's daily gasoline sales as an aggregation of purchase decisions made by individuals who are searching and learning along their predetermined travel routes. Our estimation result suggests that consumers update their beliefs of the current price distribution rather quickly with each newly observed price. This trend is consistent with the fact that the distribution of gasoline prices changes regularly, and as a result, prices sampled today are more informative about the current price distribution than past price observations. This modeling of the learning process not only allows us to estimate a more realistic and flexible formation of consumers' beliefs, but also enables us to answer the question of how much consumers would benefit if they knew the current price distribution instead of assuming it is currently the situation. After controlling for the spatial and vertical differentiation of gasoline stations, counterfactual analyses suggest that consumers would save 1.12 cents per gallon of gasoline purchased if uncertainty about the price distribution were removed, and they would save 1.60 cents per gallon if they had perfect information ex-ante about prices. Although the estimated savings are small, they are non-trivial, since gasoline demand at the station level is highly price sensitive, with an estimated elasticity on the order of -9. Moreover, it suggests that uncertainty about the actual price distribution is the primary component of imperfect information. For instance, providing consumers with information about the current price distribution alone can result in a realization of 70 percent of the total saving if they were to know all of the prices.

Our approach is also unique as it uses data on consumers' travel patterns to construct the distribution of the consumers' search sequence of gasoline stations. Most empirical papers that structurally estimate consumer search models using only market-level data assume consumers' search sequence to be random because individuals' search histories

are unobserved. However, travel patterns naturally constrain the search order when sellers have physical addresses.[5] Even for sellers without physical addresses, the order of visits can be affected by such constraints as the layout of a web-page.[6] By taking into account traffic data, we are able to improve upon the assumption of random sampling and replace it with observable variation, allowing the model to identify more realistic substitution patterns. While random sampling implies homogeneous substitution patterns between any two sellers, we find that 35 percent of the variation in the estimated elasticity of substitutions is explained by the number of common drivers shared between two stations. In other words, two stations are close competitors if they share a large number of common drivers, regardless of the distance between them. One station's price may not affect another's residual demand if they do not share any common drivers.

Our panel data allow us to estimate separately parameters that govern the consumer learning process and the distribution of the costs for postponing purchase. As volatile wholesale costs change the distribution of prices from day-to-day, variation in consumers' prior beliefs, which are based on past observed prices, is generated. The variation in consumers' prior beliefs across days leads to variation in consumers' purchase behaviors across days that is independent of the distribution of the costs for postponing purchase. The weight consumers place on their prior beliefs determines how quickly their posterior beliefs are updated with the newly observed price information. Our model identifies this weight based on how consumers' purchase behavior today is affected by the information obtained in the past.

The contributions of this paper are threefold. First, it is the first paper to estimate how quickly consumers update their prior beliefs of the unknown price distribution without consumer-level search information. Second, by doing so, we are able to quantify the additional amount of market inefficiency created by the uncertainty of the price distribution,

---

[5]Arbatskaya (2007) develops a theoretical model of ordered consumer search, finding that sellers' prices and profits are affected by the order of search.

[6]As evidenced by Kim et al. (2010), if two products sold online are not being searched together, the cross-price elasticities are numerically zero. There is also growing literature focusing on improving online platform design to reduce search frictions and increase competition (Dinerstein et al., 2018).

an outcome of consumers searching too much or too little compared to the optimum if they knew the actual price distribution. Third, this paper uses data on consumer travel patterns to simulate their search histories in a consumer search framework, thus relaxing the assumption of random sampling widely adopted by search literature when search histories are not observed. Although this feature is most relevant for the retail gasoline market, as more data about consumers' GPS movements data become available, they can be utilized to formalize how consumers search when their individual search histories have not been observed for a variety of markets.

1.2 Data

In the absence of micro-data of consumers' search histories of gasoline stations and purchase decisions, we use market-level data to make inferences about consumers' search and learning behavior that leads to gasoline purchase. Our sample consists of 47 gasoline stations in a mid-sized city, with a population of approximately 75,000 in the urbanized area.[7][8] The sample period runs from December 14, 2014, to May 31, 2016, for a total of 535 days, during which time we observe daily price of gasoline at all 47 stations and the daily gasoline transaction volume for 35 of these stations. We combine gasoline data with data on traffic flows for our estimation. Under the assumption that consumers purchase gasoline along their predetermined travel route, daily transactions at a gasoline station, as an outcome of individual purchase decisions, provide information about their searching and learning behavior concerning prices. The empirical distribution of these travel routes are constructed based on the data on traffic flows. In the following subsections, we describe the three main sources of data used for our empirical analysis.

---

[7]City name is not disclosed for the protection of the identities of the gas stations.

[8]Since the main focus of the paper is to study consumers' search behavior in the gasoline market, we excluded 14 stations from the city that do not meet the criteria of a station that competes for gasoline sales among the searching consumers. These criteria include a large price board where gasoline prices are distinctively visible and advertised to consumers, a large forecourt, and easy accessible entrance and exit. Almost all excluded gasoline stations are mom-and-pop stations with a small gasoline sales volume and primarily rely on sales inside the convenience store. Seven of eight excluded stations with matching transaction data have gasoline sales less than or at the same level as the smallest volume stations in the sample.

1.2.1 Gasoline Price Data

We collected per gallon prices of regular unleaded gasoline from two gasoline price gathering websites. The primary source was MapQuest.com, an online web mapping service whose gasoline price data are provided by Oil Price Information Service (OPIS).[9] We visited MapQuest.com daily and web-scraped the prices for every station in the city. However, not every station's price was updated daily by the website. For an average station in the city, a new price was updated on 54 percent of the days.[10] To address the issue of missing prices, we complemented MapQuest.com's data with the price data collected daily from GasBuddy.com.[11] Unlike MapQuest.com, prices on GasBuddy.com are uploaded by volunteer spotters in the area. To minimize any issues caused by the potential inaccuracy of the prices reported on GasBuddy.com, we filled in only the station-day missing MapQuest.com price with prices from GasBuddy.com.[12] The matching of stations across two data sources was based on the geometric coordinates of the stations, cross-validated with Google Map's geometric coordinates to ensure accuracy.[13] After merging the price data, a station has, on average, missing prices for 9.2 percent of the sample days. We filled in these missing prices with the most recent price observed at that station. The average time lapse for which we imputed price was 1.6 days.[14] In addition to price data, we also obtained information on station characteristics, including name, brand, address and geometric coordinates.

---

[9]OPIS obtains price information from credit card transactions and direct feeds from gas stations.

[10]The price coverage rate is slightly lower than other studies that directly use OPIS price data. One possible reason is that the city of interest is a mid-sized city with a higher share of low-volume stations than the major cities studied by other papers. Fewer credit card transactions result in fewer price feeds to OPIS.

[11]Gasoline 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).

[12]Atkinson (2008) shows that although occasional error of the published prices may occur on GasBuddy.com, these prices are still able to identify the features of retail gasoline price competition accurately.

[13]A 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.

[14]In our structural estimation, missing prices cannot be accommodated, because a station's price affects many stations' sales through the traffic network, and one missing price will result in a large number of lost observations.

1.2.2 Gasoline Transaction Data

Daily expenditure data at the station level were obtained from a major financial services provider.[15] These data reflect the total dollar amount of purchases from all debit and credit card users of this provider at each gas station on each day. We obtained two types of expenditure data from the card users for 35 of the 47 stations, pay-at-pump and in-store purchases. However, to eliminate the measurement error caused by non-gasoline transactions, we use only pay-at-pump transactions. We construct a daily measure of the total quantity of gasoline purchased at each station by dividing the total pay-at-pump expenditures by the price of regular unleaded gasoline at the station on that day.[16] We believe that our measure of the total quantity of gasoline transacted at each station provides a good representation of the behavior of consumers searching for gasoline.

1.2.3 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 prior to this station as well as the price expected for the remaining stations on the route. As a result, the search sequence of stations is essential for our analysis of consumers' searching and learning behavior in the market. In this study, we define a search route as a unique set of an ordered sequence of stations visited. In our model of consumer search with learning, an individual's decision is then modeled at the search route level. We constructed an empirical distribution of search routes describing the predicted number of consumers to drive past a specific ordered sequence of stations on an average day.

First, we constructed the empirical distribution of the travel routes comprising two elements: (i) the number of drivers traveling from an origin to a destination, and (ii) the route that drivers took along the street network connecting the two points. For the

---

[15]The name of the provider as well as the names and locations of stations in the data have been withheld to protect confidentiality.

[16]This calculation 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.

first element, we obtained the Origin and Destination Table from the state's Department of Transportation,[17] which contains information about the estimated number of drivers traveling from one Traffic Analysis Zone (TAZ) to another TAZ. This Table covers the entire county that includes the city of interest as well as six additional surrounding counties. The entire area is divided into approximately 1800 TAZs, each also associated with census information about the residents living in the zone. Table 1.1 summarizes the TAZ and the associated demographics of the residents in the county where the city of interest is located. We can see that the TAZs are extremely fine, with 75 percent of the traffic zones occupying an area of less than 1.5 km^2. The larger traffic zones are at the fringe of the county with few residents.

Table 1.1: Summary Statistics of the Traffic Analysis Zones (TAZs)

	mean	std	25%	50%	75%
Area size (km^2)	2.71	6.64	0.13	0.49	1.49
Population	668.31	920.13	141.50	357.00	733.00
Pop income > 50k	115.62	191.79	11.62	46.46	127.04

Note: The number of Traffic Analysis Zones is around 300. The exact number is not reported for the protection of the identity of the city.

To compute the route that drivers take traveling from an origin TAZ to a destination TAZ, we assume that all drivers take the travel route that minimizes driving time. We selected the centroid of the origin and destination TAZ as the drivers' start and end locations, and for every origin and destination TAZ pair, a single fastest travel route following the street network in the area was generated. We used the ArcGIS Network Analyst package to calculate the fastest travel route and the road network data set was obtained from ArcGIS StreetMap North America.[18][19]

---

[17] The Origin and Destination Table is an output of the travel demand model constructed by the Department of Transportation to forecast the traffic in year 2020.

[18] ArcGIS 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.

[19] To 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.

Figure 1.2: A Travel Route with Stations Passed
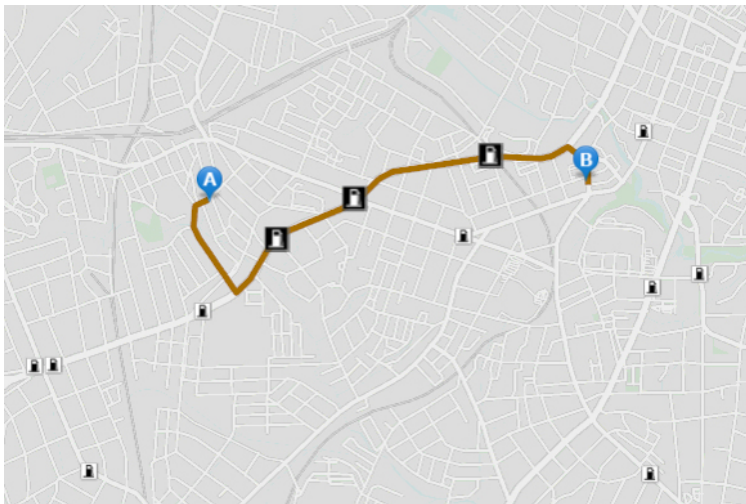
Note: This figure represents the road network of a random location.
The driving time is less than 5 minutes.

Second, we further found the identities of the stations on each travel route and the order of visits, by ranking driving distances of the stations to the origin TAZ in ascending order. Figure 1.2 provides an example of a travel route connecting a start location A and an end location B, including the three stations visited by drivers on this route. The set of an ordered sequence of stations formalizes a search route. Because difference travel routes may go through the same set of stations in the same order, we further collapse the travel routes into a search route level. We obtained the empirical distribution of search routes by summing up the predicted number of drivers driving past the same search route. We then excluded search routes with fewer than 20 drivers traveling per day. As a result, a search route abstracts from any travel time and distance and only contains two elements: (i) a unique set of ordered sequence of stations, and (ii) the number of drivers passing through it along with their demographics.[20]

Table 1.2 shows the summary statistics for the resulting travel routes and search

---

[20]Among the 47 stations, 2 stations entered the market at different times in the middle of our sample period. We generate 796 search routes before any entry, 810 search routes after the first entry, and 841 search routes after both entries. In the structural estimation, the sample periods are divided into three parts based on the date of entry and the empirical analysis is based on the empirical distribution of search routes in each period respectively.

routes. The median travel time of a travel route is 22 minutes, and a search route involves 5.3 stations and 289 drivers on average. In addition, the average percentage of drivers with an annual income of more than $50,000 is 17.98 percent on a search route.[21]

Table 1.2: Summary Statistics of Travel Route and Search Route

	count	mean	std	25%	50%	75%
	Travel Routes					
Total minutes	86099	29.80	22.99	13.00	22.12	38.42
	Search Routes					
N. stations	841	5.30	3.01	3.00	5.00	7.00
N. drivers	841	288.72	615.91	43.70	102.39	266.39
Income > 50k (%)	841	17.98	3.76	16.02	17.37	20.53

Note: Approximately 20 percent of the travel routes that have no stations are excluded.

Even though our data have many advantages, we discuss a few limitations. First, only a limited number of stations are matched with transaction data. In the estimation, although all stations' quantities transacted are predicted by the model, identification is based on the stations with quantity data observed. Due to the limited variation across stations, we can only estimate a small number of station characteristics that contribute to the vertical differentiation of stations. For this reason as well as to protect each retailer's identity, stations are broadly grouped into three categories: major-brand retailers, large-format independent retailers, and other independent retailers.[22] Second, the Origin and Destination Table describes only the travel patterns of local drivers, and it does not include information on the number of outside drivers traveling on the interstate highway, meaning a large proportion of demand at the stations near the exits of the interstate highway is not measured in our data. In addition, as these outside drivers do not observe gasoline prices along their travel route (outside drivers have to deviate from their travel route, i.e., exit the highway, to sample prices), their purchase behavior is different from the focus of

---

[21]This number is calculated using the total population whose income is greater than $50,000 divided by the total population in all origin and destination TAZs that comprise a search route.

[22]Major branded retailers are branded by the refinery companies that exclusively supply them gasoline. Large-format independent retailers have large convenience stores and many pumps. Other independent retailers are stations that do not have a large-format presence.

this paper. To account for this issue, in our estimation, we use a dummy for each of the stations located at the exit of the interstate highway to adjust for the difference between the quantity predicted in the model and the quantity observed for these stations.

1.3 Retail Gasoline Market Overview

Before introducing the structural model, it is valuable to discuss the features in the retail gasoline market that motivate our modeling choice. More specifically, we first examine the relationship between the average transaction volumes of the stations and their price reputation and location. Then, we discuss the sources that result in consumers' imperfect price information and highlight that a search with learning model is more appropriate for formalizing consumers' behavior than the standard consumer price search model in this market.

1.3.1 Relationship Between Average Station Transaction Volume and Station Characteristics

Table 1.3 describes the distribution of the station average price for all 47 stations in the city and the distribution of the station average quantity transacted for the 35 stations with matching gasoline transaction data over the sample period. The first two rows present the characteristics of all stations in the city. The average gasoline price in our sample is $1.97 per gallon, while the average volume transacted at a station is slightly below 1000 gallons per day. The interquartile range of the station average price is 10 cents per gallon, accounting for approximately 5 percent of the average price.[23] In contrast, considerable heterogeneity is found in the average quantity transacted across stations, with the station at the third quartile being six times the size of the station at the first quartile. The remaining four rows show the distribution for the large-format independent and the major-branded retailers. Comparing the distribution of the stations in each of the two categories with

---

[23]Interquartile range is calculated by the difference between the 25th percentile and the 75th percentile.

the distribution of all stations in the city, respectively, we can see that there is a negative correlation between stations' average price and average transaction volume, consistent with consumers preferring stations that tend to have lower prices. All large-format independents in our sample have average prices in the lowest quartile, while their quantity transacted is in the highest quartile. On the other hand, the average prices for all the majors are dispersed in the upper three quartiles, with their sales in the lower three quartiles.

Table 1.3: Summary Statistics of Station Average Price and Quantity Transacted

	count	mean	std	25%	50%	75%
	All Gasoline Stations					
St. avg. price ($)	47	1.97	0.07	1.93	1.94	2.03
St. avg. quantity (gl.)	35	958.87	1173.29	240.19	375.05	1428.15
	Large-format Independent Retailers					
St. avg. price ($)	5	1.93	0.01	1.93	1.93	1.93
St. avg. quantity (gl.)	5	2699.12	1677.34	2063.04	2224.84	2239.03
	Major Branded Retailers					
St. avg. price ($)	17	2.02	0.06	1.96	2.03	2.06
St. avg. quantity (gl.)	13	392.26	423.76	120.77	293.09	375.05

In addition to a negative relationship between the average price and the average transaction volume of the stations, Figure 1.3 depicts a positive correlation between the daily traffic volume and the average transaction volume at a station, both measured in logarithms.[24] As this figure suggests, as more drivers pass a station, the sampling probability of this station is higher, which, in turn, leads to higher sales. This positive relationship also confirms the advantage of using traffic data to simulate consumers' search patterns for gas stations over the assumption of random sampling with equal probability. The latter assumption, which is a simplification widely adopted by empirical literature on consumer search when individual search histories and quantity data are not observed (e.g., Hong and Shum 2006; Wildenbeest 2011; and Nishida and Remer 2018), implies that when prod-

---

[24]One outlier station is excluded from the figure. This station is in a remote area by a lake with a high sales volume, perhaps resulting from the high demand for gas for boats. This is outside the scope of the paper. The average sales volume of this station is controlled for in the structural estimation.

13

ucts are homogeneous and sampled with equal probability, the average sales volumes across stations are expected to be the same. It is also important to notice that the relationship between traffic and transaction volume is positive but not perfectly linear, indicating that additional attributes of stations also matter, such as, price reputation and brand quality, among others.

Figure 1.3: Sales Volume and Traffic Volume
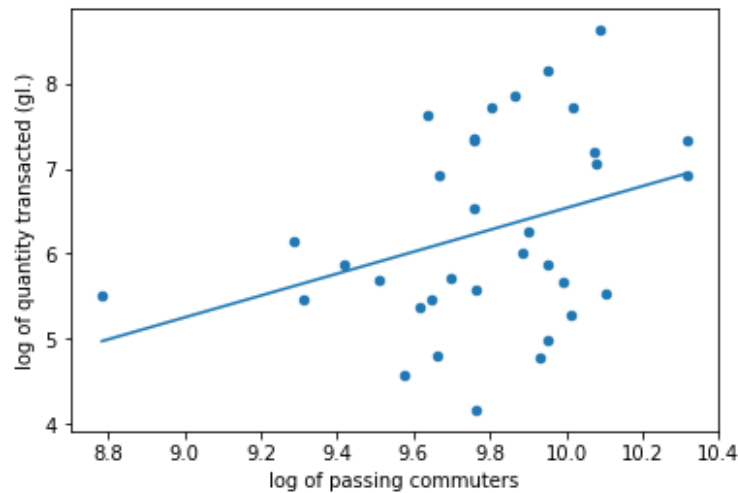
Note: The slope of the linear fitted line is 1.28, significant at 5 percent level.

Consumers also pass enough stations to allow them to search without deviating from their travel route as we will assume in the structural model. Figure 1.4 shows the distribution of the number of stations on drivers' travel routes. More than 50 percent of drivers pass at least three gasoline stations on their trips. The average number of stations passed by a driver is 3.57, a number larger than that searched by consumers before buying an MP3 player or a book from online retailers.[25]

---

[25]Using 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.

Figure 1.4: Share of Drivers By the Number of Gas Stations on Travel Route
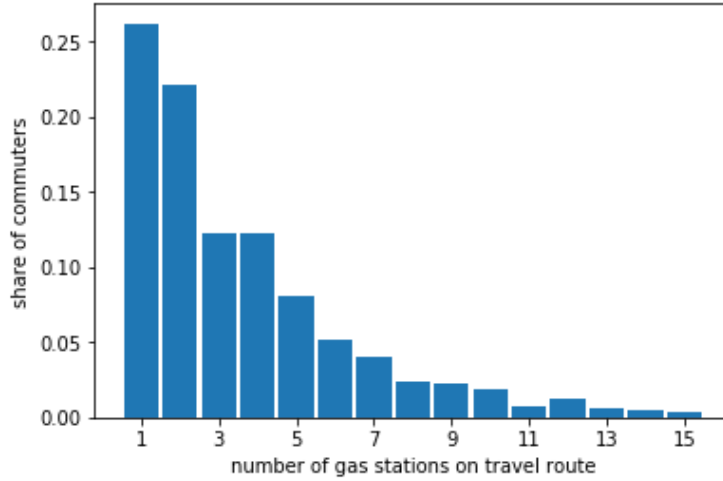
1.3.2 Two Types of Uncertainty

Frequent price changes in the retail gasoline market make it difficult for consumers to maintain accurate price information. Figure 1.5 shows that the distribution of the proportion of stations that change their price from the previous day. On an average day, about 30 percent of stations in the sample change their price at least once,[26] resulting in two types of uncertainty in the market: (i) ex-ante uncertainty about the price at each station, and (ii) uncertainty about the probability distribution of all stations' prices. To further analyze the different sources contributing to these two types of uncertainty, we perform the following basic fixed effect regression,

$$p_{j,t} = \sum_{j=1}^{J} \psi_j Station_j + \sum_{t=2}^{T} \phi_t Day_t + \nu_{j,t}, \qquad (1.1)$$

where price at station j on day t, $p_{j,t}$, is a combination of the station fixed effect, ψ_j, day-of-sample fixed effect, ϕ_t, and an idiosyncratic error term, $\nu_{j,t}$.

Not all of the price components cause uncertainty. The average price differences

---

[26]This number is only a conservative measure of the proportion of stations changing their prices on a day, due to the existence of missing prices. It is likely that price changes happen, but the data did not record them.

Figure 1.5: Proportion of Stations Changing Price From Previous Day
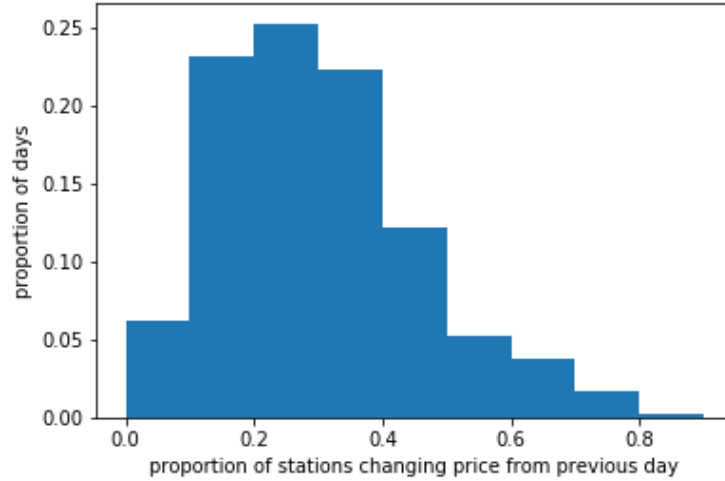
across stations, as measured by the differences in station fixed effects (ψ_j), are probably perceived by consumers. As they repeatedly observe the prices of a number of stations during their everyday driving, they tend to know the identity of higher and lower priced stations on their travel route.

After subtracting the persistent price differences across stations from the observed prices, we obtain the time variant components of the prices. We define time-variant price, $\tilde{p}_{j,t} = \phi_t + \nu_{j,t}$, as a combination of the day-of-sample fixed effect (ϕ_t), which is driven by changes in market conditions (for example, wholesale cost) common to all stations, and the price residual ($\nu_{j,t}$), which is a result of the private shock (for example, private cost shocks) specific to a certain station from day to day.[27]

These two components contribute to uncertainty in the market in different ways. First, variation in the price residual ($\nu_{j,t}$) causes station prices to fluctuate relative to one another (e.g. Lewis 2008). More specifically, station specific shocks might result in a lower priced station having a higher price than a higher priced station on a nontrivial number of days (e.g., Chandra and Tappata 2011 find that this situation occurs approximately 15 percent of the time). Consequently, price residuals create the first type of uncertainty by

---

[27]Some search literature consider price residual as a result of firms employing mixed strategy. This paper assumes pure strategy equilibria similar to Hortaçsu and Syverson (2004).

preventing consumers from knowing a station's location in the price distribution prior to search. Price residuals, which reflect the true price differences across stations that are ex-ante unknown to consumers on a day, are the focus of most of the empirical papers that structurally estimate models of consumer search. These studies typically assume that consumers know the distribution of price residuals. In particular, market conditions (reflected in ϕ_t) are either assumed to be constant over time (e.g., Hong and Shum 2006; Wildenbeest 2011; etc.), or the changes are assumed to be known to consumers (e.g., Nishida and Remer 2018 recover estimates of search costs in the gasoline market assuming that consumers know ϕ_t).

Table 1.4 summarizes the dispersion of price residuals. The size of the residual price differences across stations, which are ex-ante unknown to consumers on a single day, shows the potential gains when searching for lower prices. Daily price dispersion is calculated by the standard deviation and the range (max minus min) of the price residuals across stations within a day. The first two rows of Table 1.4 report the city-wide result. On an average day, the standard deviation and range of the price residuals are 0.066 and 0.251 dollars, respectively. However, not all of the stations in the city are easily accessible to each consumer. The dispersion of price residuals of stations along a travel route gives a better picture of the gains of searching that can be exploited by consumers driving on that route. The last two rows present the price dispersion on the search route and at the day level.[28] The average ex-ante unknown price difference across stations on a search route on a day measured in standard deviation (range) equals 0.024 (0.055) dollars. Thus, on average, an individual can realize a maximum gain of 5.5 cents per gallon of gasoline purchase by searching the stations on her travel route.

The second component of price variation and uncertainty, which results from changing market conditions as measured by ϕ_t, is a critical feature in the retail gasoline market, and it creates the second type of uncertainty. As shown by Figure 1.1, the daily distribution of prices fluctuates frequently and significantly in response to wholesale cost changes. It

---

[28]The calculation is based on search route-day that has at least two price observations.

Table 1.4: Daily Dispersion of Price Residuals ($/gal)

	mean	median	min	max	N. obs.
In City					
Standard deviation	0.032	0.032	0.013	0.069	535
Range	0.158	0.147	0.052	0.455	535
On Route					
Standard deviation	0.024	0.019	0.000	0.292	402516
Range	0.055	0.044	0.000	0.455	402516

is unlikely that consumers know the current distribution of time-variant prices with any degree of certainty.

Consumers' best guess of the distribution of prices today is likely to be based on the prices observed from past trips. Therefore, we assume in the structural model that consumers formulate their prior belief of price distribution on prices they observed while driving in the past. As consumers observe new price, they Bayesian update their prior belief based on the newly acquired price information. The update is not perfect, however, because consumers do not know why prices are changing. When the price observed at a station today is higher than the previous day, consumers do not know if this new price is specific to this station or it reflects changes in the market condition that are common to all stations. The extent to which consumers' prior beliefs of the current price distribution are updated depends on how much a price observed at a station reflects the change in the market condition. For instance, if consumers believe this price change is specific to this station, then they do not adjust their beliefs about the prices at other stations. If, on the other hand, they believe that this price change is common to all stations, then they will adjust their beliefs about the prices at the other stations. In addition, since consumers observe only a limited number of new prices on their travel route, their information about the current price distribution is incomplete. As a result, consumers' prior beliefs always impact beliefs about the current price distribution, but the weight placed on these priors decreases as more new information is obtained.

An imperfect knowledge of the price distribution can have important effects on

consumer behavior. For example, this lack of information explains the commonly observed pattern of asymmetric price response to cost changes where cost increases are passed through more quickly than cost decreases, as depicted in Figure 1.1. Lewis (2011) argues that when cost increases push up prices, consumers' beliefs of the price distribution may be lower than the actual price distribution, causing consumers to increase their search. As a result, price margins go down. On the other hand, when costs and prices fall, the current prices may be lower than consumers expected, meaning they search less, which leads to higher profit margins.

The behaviors of the equilibrium prices are consistent with the implication of a model of consumer search with learning. In the following sections, we propose a model of consumer search with learning, and we then apply the model to data and structurally estimate how consumers form their believes.

1.4 Model

1.4.1 Utility

Demand for gasoline is characterized by consumers searching for it from gas stations on their travel routes. We consider a city containing a set, $\mathbb{J}$, of J stations indexed $j = 1, 2, ..., J$. We assume that consumer each demands 10 gallons of gasoline and all share a common utility function. Consumer's indirect utility for gasoline at station j on day t is equal to

$$u_{j,t} = X_j \beta - p_{j,t},$$

where X_j represents station j's non-priced characteristics, and $p_{j,t}$ is the unit cost of gasoline (per gallon price observed multiplied by 10 gallons) at gas station j on day t. Note that the coefficient on the price is normalized to -1, so utilities are expressed in terms of dollar values.

Based on the decomposition of prices in Equation 1.1, $p_{j,t} = \psi_j + \phi_t + \nu_{j,t}$, we can

19

rewrite consumer's indirect utility as

$$u_{j,t} = X_j\beta - \gamma\psi_j - \phi_t - \nu_{j,t}$$
$$= V_j - \tilde{p}_{j,t}, \tag{1.2}$$

where ψ_j is station j's average price, ϕ_t represents the day-to-day changes in the city average price, and $\nu_{j,t}$ is the idiosyncratic deviation of station j's price on day t from its own average and the city average. We further partition utility into $V_j = X_j\beta - \gamma\psi_j$, which measures the value of station j's time invariant characteristics, and a time-varying price component $\tilde{p}_{j,t} = \phi_t + \nu_{j,t}$. We use ψ_j as a measure of a station's reputation of being a high or low priced station, and we allow consumers to respond to the differences in price reputation across stations separately from changes in time-variant prices. The difference in response is measured by $\gamma - 1$.

The partition of the utility function is motivated by the features of the retail gasoline market. Repeated observations of a number of stations and frequent purchases of gasoline at these stations allow consumers to become aware of the characteristics of stations that are constant over time, such as the location, brand, and price reputation. V_j then represents the part of utility known to consumers prior to 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.

Our approach allowing for ex-ante differentiation of products relaxes the assumption adopted by papers on search using market level data that products are ex-ante homogeneous (e.g. Hortaçsu and Syverson 2004). Kim et al. (2010) and Moraga-González et al. (2015) allow ex-ante differentiation of products, however, in their papers the sampling process is endogenous. In particular, consumers search the alternatives that yield the highest expected utilities first based on the characteristics observed. In our model, the search order of stations is exogenous as it is predetermined by consumers' travel routes.

As consumers drive along their travel route, they observe prices to realize the time-

variant price at each station. Consumers then Bayesian update their belief of the current distribution of time-variant prices with each newly observed price. In the following section, we detail the consumer learning process.

1.4.2 Consumer Learning

Let $\boldsymbol{\rho} = (\rho_1, \rho_2, ..., \rho_H)$ be a set of all possible values of time-variant prices. These values are sorted in increasing order, with ρ_1 being the lowest possible value and ρ_H being the highest. The probability of sampling each value is given by a vector $\boldsymbol{\pi} = (\pi_1, \pi_2, ..., \pi_H)$. Consumers consider these probabilities to be random variables that are distributed according to a Dirichlet distribution, following Rothschild (1974), Koulayev (2013) and De los Santos et al. (2017). We assume consumers' prior beliefs for the probabilities follows

$$f(\boldsymbol{\pi}|\boldsymbol{a}) = \frac{\Gamma(\sum_{h=1}^{H} a_h)}{\prod_{h=1}^{H} \Gamma(a_h)} \prod_{h=1}^{H} \pi_h^{a_h - 1}, \quad \sum \pi_h = 1, \quad \pi_h > 0, \tag{1.3}$$

where Γ is the gamma function and $\boldsymbol{a} = (a_1, a_2, a_H)$ is a vector of concentration parameters, which can be re-written as $\boldsymbol{a} = \alpha \boldsymbol{\omega}$, where $\boldsymbol{\omega}$ is a vector of the expected value of the probabilities of prior beliefs

$$\boldsymbol{\omega} = E(\boldsymbol{\pi}|\boldsymbol{a}) = (\frac{a_1}{\alpha}, \frac{a_2}{\alpha}, ..., \frac{a_H}{\alpha}), \quad \alpha = \sum a_h, a_h > 0, \tag{1.4}$$

and α is a scale that is often interpreted as the weight on the initial prior belief (prior sample size). The ratio a_h/α represents the expected likelihood of value ρ_h being sampled.

The prior belief is updated as consumers observe each price on their travel route. Since the Dirichlet distribution is the conjugate prior of the multinomial distribution, the posterior distribution will be Dirichlet as well. Let s_k be the set of k stations that consumers observe on their travel route. $\boldsymbol{N}_{s_k} = (n_1, n_2, ..., n_H)$ then represents the number of times each time-variant price has been sampled during the search of k stations. The posterior

expected value of the sampling probabilities is

$$E(\boldsymbol{\pi}|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}) = (\frac{a_1 + n_1}{\alpha + \sum n_h}, \frac{a_2 + n_2}{\alpha + \sum n_h}, ..., \frac{a_H + n_H}{\alpha + \sum n_h}). \tag{1.5}$$

Consumers' posterior beliefs of the price distribution depends on three aspects, the shape of the prior belief, presented by $\boldsymbol{\omega}$; its weight presented by α; and the newly obtained information, $\boldsymbol{N}_{s_k}$. According to the learning process, the expected likelihood of observing ρ_h, $\frac{a_h + n_h}{\alpha + \sum_h^H n_h}$, increases every time it is observed. The speed of consumer belief updates is inversely related to α. More specifically, the more confident consumers are in their prior beliefs, the larger the α and the slower the beliefs are updated with each new observation. As more prices are sampled from the current distribution, the posterior belief converges to the actual price distribution.

1.4.3 Ordered Search

Consumers are assumed to search the stations along their travel route. In practice, drivers commonly encounter plenty of options without having to go out of their way. The average driver in our data observes 3.57 stations on her travel route. Therefore, any deviation from the travel route, including recall (driving back to a previously passed station) is assumed to be too costly. As a result, consumers' search of gasoline stations is sequential and ordered, as they know ex-ante the predetermined order of visits of differentiated stations. In this section, we discuss how a consumer makes a decision to purchase gasoline along the route. Then in the following section, we aggregate all driver decisions to obtain the total volume of gasoline transacted at a station.

A search route s is defined as a set of sequentially ordered stations that a consumer will drive past on her trip, with $\mathbb{S}$ as the set of all search routes. At each station on the search route, a consumer updates her belief of the price distribution with each new price observed before deciding whether to purchase gasoline at this station or go to the next one. As the travel route is predetermined, observing the price of the next station is assumed to

22

be costless. The consumer maximizes her utility of a gasoline purchase by comparing the realized utility at each station with the expected value of the remaining options, which is characterized as reservation utility.

Let there be K stations, with k representing the location of consumer i on search route s, and $k = 1$ being the first station and $k = K$ the last. Let $s(k)$ return the station index j for the kth element in s, with $s(k) \in \mathbb{J}$.

Consumer i's decision at the kth station can be formalized as follows. We begin by dividing the search route into two parts, $s_k \cup \bar{s}_k$, with s_k being the set of stations already visited by consumer i, and $\bar{s}_k$ containing the remaining stations on the route. Then, based on the Equation 1.5, her belief about the current distribution of time-variant price after observing prices at the first k stations is $E(\boldsymbol{\pi}|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k})$. Conditional on her belief, she derives the reservation utility at the current kth station using backward induction from the outside option.

If she does not buy from the final station, she is allowed the alternative of waiting to purchase gasoline on a future trip. However, she needs to pay a cost of postponing purchase c_i to access the alternative. This c_i will be higher for those running out of gas or who need to purchase now for some other reasons, and lower for those who are seeking to purchase gas but not under pressure to do so immediately. Based on the observation that a driver regularly travel different routes on different occasions, the outside option is simplified as sampling a station randomly from the set of all stations, $\mathbb{J}$. Denote $Z_x(\cdot)$ as the reservation utility that represents the expected value of having x stations remaining on the route. Thus, at the last station where there are no stations remaining, the reservation utility $Z_0(\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}, \boldsymbol{V}_\emptyset, \boldsymbol{V}_\mathbb{J}, c_i)$ solves the following equation,

$$c_i = \sum_{j=1}^{J} \lambda_j \sum_{\rho_h < V_j - Z_0} (V_j - \rho_h - Z_0) E(\pi_h|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}), \tag{1.6}$$

where λ_j is the probability of sampling station j on future trips, $\boldsymbol{V}_\emptyset$ is a vector of the values of the station attributes remaining on the route, and $\boldsymbol{V}_\mathbb{J}$ is a vector of values of

station attributes outside of the route. Later in the empirical exercise, we assume equal probability of sampling a station on an individual's future trips, $\lambda_j = 1/J$. The right hand side of Equation 1.6 describes the expected gains of choosing to buy gasoline during future trips, if the utility at the last station is at Z_0. By equating expected gains with the cost of postponing purchase c_i, Z_0 summarizes the expected value of the outside option, conditional on the driver's belief at the kth station. Thus, if the realized utility at the last station is higher than Z_0, the cost of postponing purchase is higher than the expected gain from postponement, and the driver will make a purchase at the final station, and vice versa.

The calculation of reservation utility at the second to last station is then

$$Z_1(\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}, \boldsymbol{V}_{\bar{s}_{K-1}}, \boldsymbol{V}_{\mathbb{J}}, c_i) = \sum_h max\{Z_0, V_{s(K)} - \rho_h\}E(\pi_h|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}),$$

$$= Z_0 + \sum_{\rho_h < V_{s(K)} - Z_0} (V_{s(K)} - \rho_h - Z_0)E(\pi_h|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}). \quad (1.7)$$

Equation 1.7 shows that at the final station (Kth) station, consumer i will make a purchase only if the realized utility is greater than the reservation utility $u_{s(K)} = V_{s(K)} - \tilde{p}_{s(K)} > Z_0(\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}, \boldsymbol{V}_{\emptyset}, \boldsymbol{V}_{\mathbb{J}}, c_i)$, conditional on her belief at the kth station; otherwise, she will choose the outside option Z_0. Therefore, the reservation utility at the second to last station, Z_1, is higher than the reservation utility at the last station Z_0 by the amount of expected gain from search at the last station.

As consumer i continues to backward solving this recursive relationship for reservation utility for each station from the end of the route to her current location, the reservation utility at the kth station is

$$Z_{K-k}(\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}, \boldsymbol{V}_{\bar{s}_k}, \boldsymbol{V}_{\mathbb{J}}, c_i)$$

$$= Z_{K-k-1} + \sum_{\rho_h < V_{s(k+1)} - Z_{K-k-1}} (V_{s(k+1)} - \rho_h - Z_{K-k-1})E(\pi_h|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}). \quad (1.8)$$

The transition of the reservation utility at each station has the following properties.

Proposition 1: Conditional on consumer i's belief, reservation utility Z_k decreases

24

in c_i for any $k = 0, 1, ..., K - 1$.

Proposition 2: Conditional on consumer i's belief, the reservation utility decreases as the number of stations remaining on the route decreases. $Z_k \leq Z_{k+1}$ for any $k = 0, 1, ..., K - 1$.

Figure 1.6 summarizes the two propositions. The three curves in the graph depict the transition of the reservation utility as the number of stations remaining on the route changes, given three different costs of postponing purchase, assuming that the consumers' beliefs of the price distribution are uniformly distributed from 0 to 1 and stations are homogeneous with attribute value $V = 2$. First, as the postponement cost is higher, the value of purchase on future trips becomes lower, as shown by the curves when the number of stations remaining is zero. Consequently, a higher postponement cost means drivers require a lower threshold at all points along the route (*Proposition 1*). A driver who has an almost empty gas tank is more likely to accept a lower utility (higher price) at any station along her route than a driver who still has plenty of gas left in the tank. Second, as the number of stations remaining on the route goes down, the reservation utility goes down (*Proposition 2*). More specifically, when consumers come to the end of their routes, they are more likely to accept lower utilities (higher prices) as they have fewer options remaining. In addition, *Proposition 2* also suggests that, if there is a station with a high value of attributes on the route, the reservation utility is higher at any station before that station, suggesting that a driver will buy from a station only if the realized utility is higher than the expected utility at any station remaining on the route.

We further illustrate how the distribution of the cost of postponing a purchase determines the proportion of consumers who choose each option on the route. At the kth station on the route, after realization of the utility u_k and updating the belief of the price distribution $E(\boldsymbol{\pi}|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k})$, we obtain the critical cutoff of postponement cost c_k^* that equates the reservation utility at the kth station with u_k. Because Z_{K-k} is a decreasing

Figure 1.6: Transition of Reservation Utility
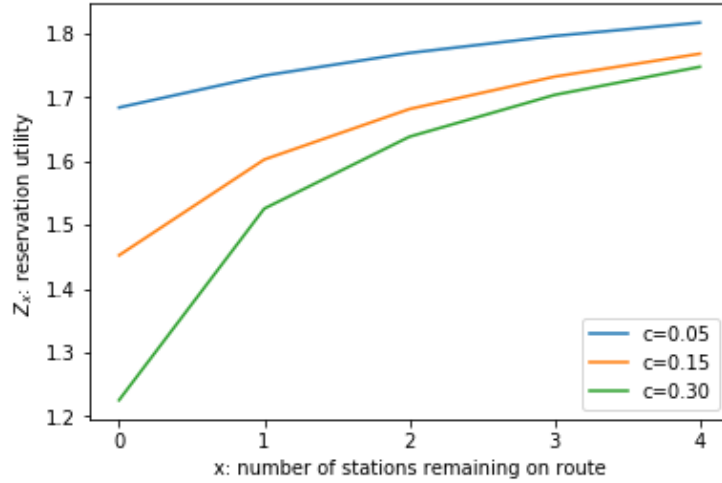
Note: This figure assumes that consumers' belief of the price distribution follows a uniform distribution with support from 0 to 1. Stations are homogeneous with attribute values equal to 2.

function in c, according to *Proposition 1*, its inverse is

$$c_k^* = Z_{K-k}^{-1}(u_k|\alpha\boldsymbol{\omega}, \boldsymbol{N}_{s_k}, \boldsymbol{V}_{\bar{s}_k}, \boldsymbol{V}_{\mathbb{J}}).$$ (1.9)

Therefore, consumers who have search cost $c > c_k^*$ will purchase at the kth station as the realized utility u_k is higher than the expected value of the remaining options, given that they have not purchased from any previously visited stations. Let G denote the distribution of the postponement costs c_i for consumers who drive on route s. Then the proportion of consumers who purchase from the kth station on the route is

$$q_k = \begin{cases} G(c_k^{**}) - G(c_k^*) & \text{if } c_k^{**} > c_k^* \\ 0 & \text{otherwise,} \end{cases}$$ (1.10)

where $c_k^{**} = min\{c_g^*$ for any $g < k\}$. As consumers with a cost of postponing purchase higher than c_k^{**} would have already bought from a station before the kth station, only when $c_k^{**} > c_k^*$, would the kth station have positive gasoline sales on that particular search route.

This concludes the formalization of the individual level purchase decision as a result of searching and leaning on each search route. In the next section, we explain the aggregation from individual to station level, and the use of daily price and quantity data at each station, combined with traffic data for estimating the model.

1.5 Estimation

Our general estimation approach is as follows. We begin by aggregating a station's share of sales, as described in Equation 1.10, across all routes the station is on to obtain the model predicted share of sales at the market level, which is the level of observation of our quantity data.

The observed market share of station, j, on day, t, $Q_{j,t}$, is defined as the share of total drivers who purchase gas at the station on that day. The number of drivers who purchase at a station is measured by the observed daily quantity of gasoline transacted at the station divided by 10 gallons, based on the assumption that a driver only buys 10 gallons of gas at each purchase, while the total number of drivers are obtained from the traffic data.[29]

Since the majority of drivers on a search route have fairly full tanks and are unlikely to even consider purchasing gas, we allow cost of postponing purchase till a future trip to have a probability mass of $1 - \eta_t$ at zero. These drivers never purchase because their expected value of future purchase, Z_0, equals the highest utility (lowest price) possible from a gasoline purchase. The remaining η_t share of drivers have a positive probability of purchasing gasoline on their search routes. Note that η_t is allowed to vary over time to account for the changes in overall demand (frequency of purchase) for gasoline. For example, during summer seasons drivers purchase more frequently, meaning the proportion of drivers with a positive probability of purchasing gasoline also goes up. We assume consumers' costs for postponing purchase follow a log-normal distribution and associate these costs to the

---

[29]A unit of gasoline purchase of 10 gallons is a scaler chosen for the convenience of interpreting the estimation results.

demographics of drivers who drive along a search route as follows,

$$lnc_{i,s} = C + \mu_c Y_S + \zeta_i^c, \tag{1.11}$$

where C is a constant, Y_S is the measures of the demographics of drivers on a route, including the proportion of drivers with incomes greater than \$50,000 dollars, and a stochastic term ζ_i^c that follows $N(0, \sigma_c^2)$ that is common to all routes. Therefore, Equation 1.10, which describes the share of drivers on search route s who purchase from the kth station, can be rewritten as

$$q_{k|s} = \begin{cases} \eta_t(\Phi(\dfrac{lnc_k^{**} - C - \mu_c Y_S}{\sigma_c}) - \Phi(\dfrac{lnc_k^* - C - \mu_c Y_S}{\sigma_c})) & \text{if } c_k^{**} > c_k^* \\ 0 & \text{otherwise,} \end{cases} \tag{1.12}$$

where $\Phi(\cdot)$ is standard normal CDF. Let $s^{-1}(j)$ returns the location index k of station j on s. The share of total drivers who buy from station j at time t is,

$$\hat{Q}_{j,t}(\boldsymbol{\theta}) = \frac{1}{\sum_{s \in \mathbb{S}} W_S} \sum_{s \in \mathbb{S}} q_{s^{-1}(j)|s,t} W_S \mathbb{1}(j \in s), \tag{1.13}$$

where W_s is the number of drivers on route s obtained from the traffic data. The set of parameters to be estimated is given by $\boldsymbol{\theta} = (\alpha, \beta, \gamma, C, \mu_c, \sigma_c, \eta_t)$, where α is the weight on consumers' prior beliefs, β is the value of station attributes such as brand, γ is the sensitivity to price reputation, and C, μ_c and σ_c determine the shape of the postponement cost distribution.

Next, we assume that the model predicted market share differs from its observed value at each station by a additive independent normal error, $\epsilon_{j,t} = Q_{j,t} - \hat{Q}_{j,t}(\boldsymbol{\theta})$, with $\epsilon_{j,t} \sim N(0, \sigma)$, and subsequently construct the likelihood function to be maximized as follows,

$$\prod_t \prod_j \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(Q_{j,t} - \hat{Q}_{j,t}(\boldsymbol{\theta}))^2}{2\sigma^2}}. \tag{1.14}$$

The error term $\epsilon_{j,t}$ represents an aggregate level measurement (prediction) error, similar to Bresnahan (1987), Hortaçsu and Syverson (2004), Kim et al. (2010), etc. Several sources can lead to measurement error. First, inaccuracy of the gasoline prices and the imputation of missing prices can directly cause measurement error in the observed market share at a station since the quantity of gasoline transacted at a station is calculated by dividing the total value of transactions by the price. Second, the model predicted market shares are based on the estimated daily average traffic flows that do not reflect temporary road repair or congestion. True daily market shares may be influenced by such temporary changes.

Consumers' prior beliefs about the price distribution on day t before observing any prices on their search route is assumed to be a result of continuous learning of the price distribution based on the observed prices on past trips, with older price observations discounted more heavily. Consumers are assumed to have a homogeneous prior belief and updating process, described as follows,

$$\boldsymbol{\omega_t} = E(\boldsymbol{\pi}|\alpha\boldsymbol{\omega}_{t-1}, \boldsymbol{N}_{t-1,\mathbb{J}}/J), \tag{1.15}$$

where $\boldsymbol{\omega_{t-1}}$ is consumers' prior beliefs on day $t-1$, and $\boldsymbol{N}_{t-1,\mathbb{J}}/J$ is the vector of the number of times each price is being observed in the city on day $t-1$ divided by the total number of prices observed, which is also the empirical density distribution of the city price distribution. One can think of this process as an exponential moving average of the city price distribution. As discussed in Section 1.3, because consumer information about price distribution on any given day is imperfect and incomplete, the prices observed in the past always carry some weight in the consumers' beliefs. This weight, measured by α, governs the speed of the learning process.

Identification of the different parameters in the model relies on different sources of variation. First, consumer preferences, β and γ, for product attributes, including brand and price reputation, are identified by observed variation in average market shares across

stations. Second, the parameters, C, μ_c and σ_c that govern the distribution of the costs of postponing purchase are identified by (i) how observed variation in market share is related to the variation in station average utilities (normalized to dollars) across stations, conditional on the variation in the number of drivers who go past each station and (ii) how variation in a station's market share over time relates to changes in day-to-day prices. More specifically, the first variation shows how likely it is for drivers assigned to a route with bad stations (those with high price reputations) to postpone purchase until they are assigned a route with better options on average, while the second variation represents intertemporal substitution reflected in a station's market share. Third, the weight on the prior belief α is identified by the variation of the difference between consumers' prior beliefs and the current price distribution over time.

1.6 Results

1.6.1 Parameter Estimates

The estimated parameters are shown in Table 1.5. The parameter estimate of the weight on the prior belief of the price distribution is 5.31. More specifically, consumers' initial prior belief, which is formulated by past price observations during everyday driving, is as informative about the current price distribution as if consumers had already observed approximately five prices today. When less weight is placed on prior beliefs, consumers update their beliefs more quickly with each new price observation. For example, if an individual believes that the average price today is $2.00, a prior weight of 5 suggests that if she observes a price of $2.12, she will update her belief of the average price to $2.02. However, if the prior weight is 11, her posterior belief of the average price is $2.01. It is important to note that consumers do not update their belief of the expected price to be equal to the price observed at a particular station because they do not know whether the change is specific to this station or is common to all stations. The estimated rate of updating is much faster than what is generally assumed by literature. In the two empirical

papers that apply models of search with learning (Koulayev 2013 and De los Santos et al. 2017), the weight on the prior belief is not estimated but rather chosen by the researchers to be the number of product-retailer combinations. This results in a much higher weight on the prior. (For example, Koulayev 2013 set α to be ranging from 24 to 82). Our estimate of a relatively fast learning process is consistent with the fact that the distribution of gasoline prices changes regularly in response to wholesale cost volatility. As a result, past price observations carry only limited information, while prices sampled today are much more informative about the current distribution of prices.

Table 1.5: Estimation Results

	Coeff.	Std. Err.
Learning		
Prior Weight (α)	5.3055	(1.4703)
Seller Attributes		
Price Reputation (γ)	-1.7872	(0.3093)
Majors (β)	-0.2419	(0.1749)
Large-format Independents (β)	2.1853	(0.1429)
Log of Postponement Cost		
Constant (C)	-2.7456	(0.1936)
Income > 50k (μ_c)	0.0773	(0.0801)
SD (σ_c)	4.8452	(0.2194)

Note: Bootstrapped standard errors are in parentheses. The coefficient estimates are for an one time purchase of 10 gallons of gasoline. Day of sample fixed effects and month of sample fixed effects are included in η_t to capture the change in the daily overall demand for gasoline. Estimates of η_t are not reported.

The estimated value of station attributes, which are ex-ante known to consumers, are compared to their response to day-to-day price changes, which are unknown prior to search. More specifically, the coefficient on the known and the persistent price reputation across stations is estimated to be -1.79, a value significantly different from -1, the response to time-variant price changes. This suggests that when consumers are driving, they are likely to postpone gasoline purchase until they reach a station that tends to be less expensive even if it means forgoing lower prices offered by a station that tends to be more expensive. If a

station's price goes up by two cents (approximately 1 percent of average gasoline price, $1.97, in the sample) in a day, which is ex-ante unknown to consumers and only revealed to those who drive by this station, it would see an average decrease in market share of 9.49 percent. On the other hand, if a station increases its persistent price reputation by two cents, which is known to all consumers ex-ante, an average station would experience a 24.65 percent drop in its market share. Table 1.6 summaries the results. This large difference in the change in market share is driven by the fact that consumers are almost twice as responsive to a change in price reputation than to a change in time-variant price. In addition, change in price reputation results in an ax-ante change in consumers' purchase probability at the station, while a time-variant change results only in a change in the purchase probability ex-post.

Table 1.6: Average Percentage Change in a Station's Market Share Under Two Types of Price Change

	mean	std	25%	50%	75%
Transitory Price (%)	-9.49	7.41	-11.56	-7.78	-4.85
Price Reputation (%)	-24.65	17.00	-31.79	-22.20	-11.07

Note: The stations near the interstate highway and the lake are excluded from the calculation.

The coefficient on the major-branded retailers is negative, however not significantly different from zero, which is the mean value of the other independent retailers. This result does not provide evidence that consumers are willing to pay a premium for branded gasoline. However, consumers are willing to pay about $2.19 premium for 10 gallons of gas at the large-format independent retailers, probably because they have standardized and consistent amenities, with large islands, plenty of pumps, and large, clean convenience stores, as well as a brand image of having a low price, in addition to their individual station price reputations.

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 6.4 cents per 10 gallons purchased, and 25 percent of these drivers have postponement costs higher than $2.09. Figure 1.7 depicts the estimated postponement cost CDF. These num-

bers indicate that the majority of drivers start seeking gasoline well in advance of their tank being empty. As they do not act as if they are pressured to buy gas immediately, they can postpone their purchase until they drive on a route with low priced stations. However, approximately 25 percent of the drivers seeking gasoline are willing to pay 20 cents per gallon more to purchase gasoline on the current route, due, for example, to an almost empty gas tank. Although the coefficient on income is positive, it is not statistically nor economically different from zero. One would expect stations on routes with a higher proportion of high-income drivers to have higher sales volumes even if prices are higher because such drivers are less price sensitive and have higher postponement costs. However, this expectation is not supported by the data, perhaps because the variation in the proportion of residents with annual income higher than $50,000 across traffic zones is too small for identification (the standard deviation is 3.76 percentage points).

Figure 1.7: CDF of the Cost of Postponing Purchase
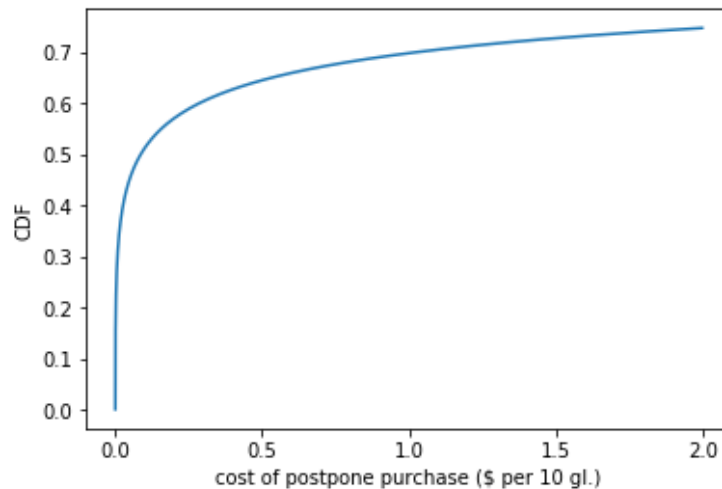
1.6.2 Cross-Price Elasticities

In this subsection, we evaluate the importance of spatial differentiation in explaining substitution patterns across stations by analyzing the impact of the distance and the number of common drivers shared between stations on the estimated elasticity of substitution. Based

on the parameter estimates shown in Table 1.5, we obtain the estimated cross-price elasticity for a pair of stations by calculating the percentage change in target station i's market share due to an ex-ante unknown one percent increase in the price of the other station, j, $\varepsilon_{i,j} = \%\Delta\hat{Q}_i/\%\Delta p_j$. Given that a majority of station pairs have virtually zero elasticities (it is not surprising that two stations located at the opposite ends of the city do not compete closely against each other as they do not share much common traffic), this analysis focuses only on pairs of stations with a cross-price elasticity greater than 0.01, a value exhibited by less than 20 percent of the total possible station pairs.

Table 1.7 summaries the distribution of the estimated cross-price elasticities and two measures of spatial differentiation between stations. We can see that even after the exclusion of station pairs that do not compete closely against each other, the distribution of cross-price elasticities is right-skewed. The average cross elasticity is 1.73, with only 10 percent of the pairs having a cross elasticity greater than 2.16. In our restricted sample, the average common traffic share between two stations, which is defined as the proportion of the target station i's traffic that also goes past station j, is 24 percent, and the average driving distance is 3.36 miles.

Table 1.7: Summary Statistics of Elasticity of Substitution and Measures of Spatial Differentiation Between Stations

	count	mean	std	min	10%	50%	90%	max
Cross Price Elasticity	284.0	1.73	8.85	0.01	0.02	0.15	2.16	127.91
Share of Common Traffic	284.0	0.24	0.28	0.00	0.01	0.11	0.72	1.00
Driving Distance (*miles*)	284.0	3.36	2.08	0.05	0.91	3.11	6.28	10.51

Note: This table includes only station pairs with a cross-price elasticity greater than 0.01, a value exhibited by fewer than 20 percent of the total station pairs.

Table 1.8 reports the results of a regression of the logarithm of the estimated cross-price elasticity on the measures of the spatial differentiation of station pairs. As expected, these regression results indicate that cross-price elasticity between two stations increases in the proportion of common traffic and declines as the driving distance between them increases. A comparison of the results in Columns (1) and (2) suggest that the proportion of

Table 1.8: Regression Results of Estimated Elasticity of Substitution on Measures of Spatial Differentiation Between Stations

| | $ln(\hat{\varepsilon}_{i,j})$ | | |
	(1)	(2)	(3)
Intercept	-0.6585**	-0.0166	0.0546
	(0.2451)	(0.2237)	(0.2426)
$ln(Distance)_{ij}$	-1.0342***		-0.3803**
	(0.1585)		(0.1419)
$ln(CommonTraffic)_{ij}$		0.6977***	0.5730***
		(0.0667)	(0.0647)
N	284	284	284
R2	0.22	0.35	0.37

Note: Standard errors, shown in parentheses, are clustered at the target station level. $CommonTraffic_{ij}$ measures the proportion of station i's traffic that also goes past station j. ***, **, and * represent significance at the 1%, 5% and 10% level respectively.

common traffic explains a more significant fraction of the variation in substitution patterns between stations (35 percent) than driving distance. Based on the coefficient estimates in Column (3), holding the driving distance between the two stations constant, a doubling of the common traffic share will increase the elasticity of substitution by 57 percent.

1.7 Counterfactual

In this section, we study the value of price information to consumers using our estimated demand system. More specifically, we calculate the net change in consumers' surplus after stepwise removal of two types of information uncertainty, (i) uncertainty about the price distribution, and (ii) uncertainty about the price at each station. To do so, we simulate consumers' purchase decisions and compute the aggregated value of these choices across hypothetical cases with different amounts of information uncertainty, while holding the supply side constant.

We begin our analysis with the current state in the presence of both types of uncertainty. The total value for all consumers seeking gasoline on day t under dual uncertainty,

$TV_{du,t}$ is calculated as,

$$TV_{du,t} = \sum_{s \in \mathbb{S}} \eta_t W_s \sum_{k \in s} u_{k,t} q_k(\alpha \boldsymbol{\omega}_t, \boldsymbol{N}_{s_k,t}) + \sum_{s \in \mathbb{S}} \eta_t W_s \sum_{k \in s} \int_0^{c_{du,s}^{**}} Z_0(c|f_t)g(c)dc. \qquad (1.16)$$

$\sum_{s \in \mathbb{S}} \eta_t W_s \sum_{k \in s} u_{k,t} q_k(\alpha \boldsymbol{\omega}_t, \boldsymbol{N}_{s_k,t})$ measures the total realized utility for consumers who have made the purchase decision conditional on their beliefs at each station on each search route, and $\sum_{s \in \mathbb{S}} \eta_t W_s \sum_{k \in s} \int_0^{c_s^{**}} Z_0(c|f_t)g(c)dc$ computes the total expected value for consumers who have chosen to postpone their purchases to future trips, where $Z_0(c|f_t)$ is the value of the outside option evaluated at the current price distribution measured by the empirical price distribution, f_t. $c_{du,st}^{**} = min\{c_{du,k}^*$ for all $k \in s\}$ is the critical cut-off postponement cost on search route s under both types of uncertainty, with consumers whose $c > c_{du,st}^{**}$ choosing to purchase on the route, and consumers whose $c \leq c_{du,st}^{**}$ choosing to postpone purchase to future trips, and $c_{du,k}^*$, solved by Equation 1.9, is the critical cut-off postponement cost for each station on the route, conditional on consumers' beliefs.

The resulting inefficiencies that occur to consumers under both types of uncertainty are illustrated by the following situations. First, when prices decline, the observed prices are likely to be lower than consumers' beliefs about them. As a result, some consumers may choose to purchase gas immediately at an early station on a route, forgoing options later in the route or future trips that may generate larger savings. Second, when prices increase, the observed prices are likely to be higher than consumer's beliefs of prices. As a result, some consumers may choose to continue searching for less expensive stations or postpone their purchases, eventually having to settle down on more expensive options before learning that the price increases are global.

To evaluate the magnitude of these inefficiencies caused by information uncertainty, we calculate the total value for all consumers under the counterfactual scenario where there is only price uncertainty and the consumers know the distribution of current prices. This

assumption is the one typically adopted in the literature on consumer search,

$$TV_{pu,t} = \sum_{s\in\mathbb{S}} \eta_t W_s \sum_{k\in s} u_{k,t} q_k(f_t) + \sum_{s\in\mathbb{S}} \eta_t W_s \sum_{k\in s} \int_0^{c_{pu,s}^{**}} Z_0(c|f_t)g(c)dc, \qquad (1.17)$$

where $c_{pu,s}^{**}$ is the critical cut-off postponement cost for consumers who are indifferent to purchasing on the current route or postponing their purchases.

Finally, we compute the total value for all consumers under the counterfactual scenario where they know the current prices at all stations. Consequently, consumers choose whether to buy from the station that provides the highest utility on a route or to postpone their purchases to a future trip. The total value under perfect information is,

$$TV_{pi,t} = \sum_{s\in\mathbb{S}} \eta_t W_s \sum_{k\in s} u_{st}^*(1 - G(c_{pi,s}^{**})) + \sum_{s\in\mathbb{S}} \eta_t W_s \sum_{k\in s} \int_0^{c_{pi,s}^{**}} Z_0(c|f_t)g(c)dc, \qquad (1.18)$$

where $u_{st}^* = max\{u_{kt}$ for all $k \in s\}$ and $c_{pi,s}^{**}$ is the critical cut-off postponement cost for consumers who are indifferent between purchasing from the station with highest utility on the route or postponing their purchases.

Based on our computation, we find that if consumers know about the current distribution of prices, their net surplus increases by an average of \$210.30 on a day in the city, calculated by $\frac{1}{T}\sum_{t=1}^{T}(TV_{pu,t} - TV_{du,t})$. This increased benefit is approximately 0.56 percent of the total value of transactions on a day in the city predicted by the model. It is equivalent to an average of 1.12 cents in savings per gallon when consumers are fully informed about the current price distribution.

Under the second counterfactual scenario, where consumers have perfect information about today's price offered by each station, their net surplus would increase by an average of \$299.90, which is equivalent to an average of 1.60 cents in savings per gallon. This estimated benefit of reduced information uncertainty seems to be small compared to the observed price dispersion in the city (summarized in Table 1.4) because our estimated model takes into account the limited number of stations on consumers' travel routes and the costs of accessing more stations not on the current travel route. More specifically, the

costs of postponing the purchase to future trips, independent of the amount of information consumers have, prevent all consumers from purchasing at the cheapest station in the city, a situation that would result in significant benefits of having perfect information. In addition, the estimated benefit of information is net of the ex-ante known attributes of the stations, including price reputation and brand qualities. As a result, our simulation predicts 90.42 percent of consumers who purchased gasoline on a route have done so from the station that generates the highest utility, while 82.69 percent of consumers have bought from the station with the lowest price under both types of uncertainty.

1.8   Conclusion

In this paper, we developed a model to study consumer behavior for searching for the best offer while learning about the distribution of prices in the context of the retail gasoline market. Using a novel panel dataset of daily station-level gasoline sales and prices, combined with data on the empirical distribution of drivers, we estimated the parameters that govern consumers' learning, utility, and cost of substitution from current purchase to future purchases. We are the first to estimate the learning process in a consumer search framework. We found that consumers place a relatively high weight on newly observed prices when formulating their beliefs of the overall price distribution, a founding consistent with the observation that prices sampled today are more informative about the current price distribution than past price observations in a market of volatile prices.

Our modeling framework has many important advantages. First, we model consumers' search sequences of gasoline stations based on their travel patterns. Consequently, search sequences of stations are simulated based on the empirical distribution of traffic flows, allowing us to estimate consumer search and learning behavior using only aggregated station-level data. In addition, the model generates realistic substitution patterns that depend on the amount of common traffic shared between two stations. Second, the model allows for ax-ante vertical differentiation of stations and is not affected by dimensionality

when integrating out individuals' search histories. We find that consumers respond to the ex-ante known differentiation of stations, including brand and price reputation by more than the actual dollar amount of the price difference. Although we do not explicitly model risk aversion in our framework, this result likely captures the value of a reduction in uncertainty, which a station or a brand provides to consumers through a consistently lower price. Third, we allow for uncertainty in consumers' beliefs about the current price distribution. Our counterfactual results suggest that information on the actual price distribution benefits consumers by an amount that is equivalent to a 1.12 cents in savings per gallon of gasoline purchased. On the other hand, perfect consumer information ex-ante about the price at each station generates an average of 1.60 cents of savings per gallon. Although the suggested benefit of reducing information uncertainty in the retail gasoline market seems small, it is non-trivial. It is also important to note that this result is obtained after controlling for the cost of substitution across spatially-differentiated stations, as well as the known average price difference across stations. Analysis that fails to take into account these two factors, for example if everyone buys from the lowest priced station in the city, would overestimate the benefit of increasing the amount of information.

There are several avenues for further development of our model, the most important being to model the supply side decision. By including the pricing decision of the stations, we would be able to analyze the amount of equilibrium price dispersion in the retail gasoline market caused by information uncertainty and spatial differentiation. In addition, our modeling framework can be applied to multiple cities and be used to empirically analyze the link between the speed of learning and the degree of asymmetric price response.

Chapter 2

Bidding Behavior in A + B Procurement Auctions: A Laboratory Analysis of the Effect of Complexity and Competition (with Bernardo F. Quiroga and Brent B. Moritz)

2.1 Introduction

A+B auctions (or scoring auctions) are widely applied to situations where attributes other than price matter. In these auctions, bidders not only bid on price (the A dimension), but also on non-monetary attributes/quality (the B dimension). Each component of a bid is evaluated based on a scoring rule, and the bidder achieving the best score wins the auction. The non-monetary dimension usually includes delivery time and any other mea-

surable quality attributes that endogenously impact the outcome of an auction as assessed by the scoring design. These auctions are common in public procurement. For example, the Department of Transportation often adopts A+B auctions to give contractors bidding on highway construction projects an explicit incentive for accelerated completion when the cost of delay for commuters is substantial (e.g. Lewis and Bajari 2011). The Department of Defense depends on scoring auctions to procure weapons systems where the non-price attributes (quality of the goods) are critical (e.g. Che 1993). The use of scoring auctions is widespread in the private sector as well. Online platforms, for instance, select the ads with the best price bid and the most relevant content (quality) to display to their users. A scoring algorithm that aims to increase both advertisement revenue and user experience is employed to determine which advertiser wins placement on the platform.

Several researchers have studied the theoretical properties of A+B auctions (e.g. Che 1993; Branco 1997; Asker and Cantillon 2008; Lewis and Bajari 2011; Gupta et al. 2015). Asker and Cantillon (2008) found that the scoring mechanism dominates other commonly used procedures for bidding differentiated products from a theoretical perspective, one example being the price-only auction with a minimum quality threshold. However, when evaluating the performance of A+B auctions in practice, researchers might be concerned about the bidders' ability to submit bids that maximize their expected utility. Experimental research by Cox et al. (1985, 1988) and many others has documented the difficulties that bidders face in making optimal offers in sealed-bid-first-price auctions, even with simplifying assumptions and price-only decisions. In the A+B auction setting, bidders not only have to decide on the price to charge but they also need to choose the quality level to fulfill. It is unclear if bidders are able to follow the optimal strategy suggested by standard auction theory. The increased complexity of the procedure may introduce noise in the bidders' decision making and, in turn, limit the mechanism's ability to realize the attractive properties, such as the efficient allocation of resources. As documented by Quiroga et al. (2019), suppliers in the laboratory show a substantial discrepancy between the theoretical prediction resulting from standard theory in A+B auctions and the observed behavior of bidders.

In this paper, we conduct an experiment similar to the A+B procurement auctions adopted by government agencies in order to empirically study the performance of A+B auction design. In this experimental setting, we examine how bidders react to the increased complexity of the A+B mechanism and how their deviation from the Nash equilibrium strategy impacts the outcome of the auctions and consequently, the performance of the mechanism. In the field, it is usually difficult to investigate the bidder behavior without imposing a series of strong rationality assumptions because of endogeneity (omitted variables) and the private cost information usually unobserved by the econometrician. The advantages of our experiment compared to field data are that bidders behave in a controlled environment and that we are able to observe bidders' private information which determines the cost structure.

Our experiment of A+B auctions consists of four treatments comprised of varying numbers of bidders and bid dimensions. This variation in competition and complexity, as we will show later in the paper, is important for identification. More specifically, two treatments involve four bidders in each auction, and in the remaining two, the number of bidders is two. In terms of complexity, bidders have to choose both a price and quality bid for two of the treatments, while they need to submit only a price bid in the remaining treatments as the optimal quality level is provided to them.

Our analysis consists of two parts. In the first part, we treat the private information (values) about the cost structure as data and study the systematic deviation of bidder behaviors from the Bayesian Nash equilibrium strategy based on the assigned private information. We investigate whether these observed discrepancies in human behavior are consistent with the predictions of competing behavioral models, with a goal of obtaining a better understanding of behavioral biases. In the second part of our analysis, we incorporate the behavioral aspects obtained from the initial analysis in structural auction models and we evaluate these models' ability to estimate bidders' private cost information. This exercise is particularly important because structural estimation methods have become increasingly popular in empirical research on auctions. In these studies, economists estimate

bidders' private values, subsequently using them to study welfare or predict bidder behaviors in different counterfactual simulations, such as changes in auction design. Therefore, the ability of structural models to generate accurate estimates of private values is essential. Throughout our structural estimation, we treat each bidder's private values as unknown in order to mirror the information accessible to econometricians using field data. To assess the performance of alternative structural models, we then measure the dfference between the estimated and the assigned private values.

The rest of the paper is structured as follows. Section 2.2 introduces the alternative models for A+B auction, followed by Section 2.3, which describes the experiment design to uncover the behavior of the bidders. Section 2.4 examines if the bidder behaviors are consistent with the models' predictions, and Section 2.5 applies structural econometric methods based on the models to estimate the bidders' private information, while Section 2.6 provides a concluding summary of this study.

2.2 The Model

This section provides the formulation of the problem under study here. We follow the notation in Quiroga et al. (2019). Consider a buyer who seeks to procure an indivisible good (or service) for which there are N potential suppliers. Suppliers, indexed by j, know the number of competitors they face, N–1, before submitting a quote or proposal. The good is characterized by its price, p, and a quality attribute, q, that can be contractually enforced.

Contracts are assigned using an A+B auction. As explained in the introduction, an A+B auction scores contractors on both price p and quality q, awarding the contract to the supplier with the highest valued score based on the scoring rule

$$S(p, q) = aq - p, \tag{2.1}$$

where a is a common-knowledge constant and positive factor and its weight, a, reflects how

much buyer value an additional increase in quality.

Supplier j's profit π_j from selling a good (p,q) (or equivalently, conditional on winning the contract) is the quoted price minus the production cost, p–$C(z_j,q)$. Here z_j is a random signal for supplier j, drawn from bounded support $[L,U]$. z_j measuring supplier j's private information about his cost structure. We model supplier cost as the sum of a quality-independent baseline cost and a quality-dependent cost structure:

$$C(z_j,q) = F + \frac{q^{1+\frac{1}{\eta}}}{(1+\frac{1}{\eta})z_j}. \tag{2.2}$$

The baseline cost parameter $F \geq 0$ is assumed to be a deterministic, non-sunk, fixed cost to be incurred by the supplier only if the contract is won and the good is sold. For simplicity, it is identical and known by all suppliers. This baseline cost can be interpreted as the cost of completing a project at the minimally acceptable quality level. For example, the quality dimension in Lewis and Bajari (2011) is acceleration time (completion time savings). The baseline cost is, therefore, interpreted as completing the project by the deadline set by original design engineer's schedule. We made this baseline fixed cost common to all suppliers when they win an auction, thus avoiding ex-ante asymmetry among suppliers as neither supplier favoritism nor skill in estimating production costs are focal points in this study.

In terms of the quality-dependent portion of the cost structure, this specification yields a marginal cost (i.e., supply) function for quality dimension given by $C_q(z_j,q) = \partial C/\partial q = q^{\frac{1}{\eta}}/z_j$, implying a constant elasticity of supply for the quality, $\partial lnq/\partial lnC_q = \eta$. Since the product quality literature in operations management usually assumes costs to be quadratic in quality (e.g., Karmarkar and Pitbladdo 1997, Karaer et al. 2015), thus implying a unit elasticity of supply for quality, we parameterized $\eta = 1$. Therefore, the marginal cost of increasing the quality level for supplier j is q/z_j, meaning the higher the z_j, the more efficient supplier j is at quality provision.

2.2.1 Bayes-Nash Equilibrium Model

2.2.1.1 Risk Neutral

If bidder j is risk neutral, then he submits sealed bid (p, q) to maximize his expected utility,

$$\max_{p,q}(p - F - \frac{q^2}{2z_j})x_j, \text{ subject to } s = aq - p, \tag{2.3}$$

where x_j is the probability j wins the auction.

A rational bidder of type z_j with score s_j will choose (p, q) to maximize profit $p - C(z_j, q)$ conditional on winning. Substituting the score constraint in the objective function yields

$$\max_{q}\{aq - F - \frac{q^2}{2z_j} - s_j\}. \tag{2.4}$$

From Equation 2.4, the optimal q does not depend on s_j. Solving for Equation 2.4, the optimal quality provision is

$$q_j^* = az_j. \tag{2.5}$$

The optimal quality equates marginal benefit with marginal cost of quality. Following Che (1993); Asker and Cantillon (2008), we define a value function, referred to as a pseudo-type, $\nu(z_j)$, measuring the maximum social welfare a bidder of type z_j can generate;

$$\nu(z_j) = \frac{a^2 z_j}{2} - F. \tag{2.6}$$

With this information, bidder j participating in this auction has an expected profit given by

$$E\Pi_j = (\nu(z_j) - s_j)x_j = (\frac{a^2 z_j}{2} - F - s_j)x_j \tag{2.7}$$

This expression is mathematically equivalent to a sealed-bid first-price auction with

pseudo-types functioning as private valuations and scores assuming the role of prices. There-fore, bidders' preferences over outcomes (x_j, s_j) are fully captured by $\nu(z_j)$. As Asker and Cantillon (2008) show, pseudo-types are sufficient statistics for this auction mechanism because once the distribution of pseudo-types is obtained, the result is a complete charac-terization of the score auction and its outcome.

Now we can solve for the optimal price portion of the bid. Let $s(\nu)$ be a monotone bidding function mapping pseudo-types ν onto score s in equilibrium. Using a result from Riley and Samuelson (1981), the symmetric Bayesian-Nash equilibrium (SBNE) score bid $b(\nu_j)$ will satisfy the following expression at optimality,

$$s(\nu_j) = \nu_j - \frac{1}{N-1}\frac{F_\nu(\nu_j)}{f_\nu(\nu_j)s'(\nu_j)}. \tag{2.8}$$

This expression fully characterizes the optimal bid strategy function, $s(\nu_j)$. This strategy will depend on the resulting distribution of pseudo-types induced by the bidders' efficiency types, as well as on the number of bidders. This characterization describes optimal bids not only at the pseudo-type and score level (given by Equation 2.8) but also by quality level (given by Equation 2.5) and price (given by difference using the scoring rule). Therefore, the optimal score and price bids are

$$s_j^* = b(\nu_j) = \nu_j - \frac{1}{N-1}\frac{G(s_j^*)}{g(s_j^*)} \tag{2.9}$$

$$p_j^* = aq_j^* - s_j^*, \tag{2.10}$$

where $G(s)$ and $g(s)$ are the CDF and PDF of the score distribution the human bidder faces.

2.2.1.2 Risk Aversion

It has been widely observed that human bidders are likely to bid higher than the Symmetric Bayes Nash equilibrium bid in first price sealed-bid auctions. One common

explanation for overbidding is that bidders are risk averse (e.g. Cox et al. 1985, 1988). Here, we model the bidder utility function as constant relative risk aversion (CRRA), $U_j(\pi) = \pi^{\theta_j}$ and extend the method proposed by Campo et al. (2011) for analyzing risk aversion for first-price auction to A+B auction; θ_j is bidder j's coefficient of relative risk aversion, with $\theta_j = 1$ representing risk neutrality.

It is easy to show that with risk aversion, conditional on winning, the quality optimum score bid and pseudo-type still follow Equation 2.5 and 2.6 respectively. Therefore, based on CRRA, the optimal score bid is presented by

$$s_j^* = b(\nu_j) = \nu_j - \frac{\theta_j}{N-1}\frac{G(s_j^*)}{g(s_j^*)} \tag{2.11}$$

and optimal price bid follows Equation 2.10. If bidders are risk averse, then $\theta \in (0,1)$. We will see that the Nash equilibrium score bid under risk aversion, $s_{RA}*$, is larger than the risk neutral Nash bid, $s_{RN}*$, conditional on private information. Intuitively, when facing uncertainty, risk averse bidders are likely to bid more aggressively to increase the probability of winning at the cost of a smaller profit if they win. The Nash equilibrium score bid still remains linear in bidders' private information.

2.2.2 Quantal Response Equilibrium Model

Because solving this differential equation is complex, we expect that bidders probably do not optimize and submit bids different from those predicted by the Bayes-Nash equilibrium. However, recent research in the experimental social sciences has used McKelvey and Palfrey (1995)'s Quantal Response Equilibrium (QRE) model to model noise in bidders' decision making and, thus, reconcile deviations from Nash equilibrium predictions. A frequently used variant of QRE is the logit model. Analogous to other discrete choice models, in the logit equilibrium a bidder's payoff is determined by his utility plus a random perturbation following an i.i.d. extreme value process.

Goeree et al. (2002) applied the QRE framework to auction settings and found that QRE model, combined with risk aversion, predicts bidder behaviors in the first price auction very well. Here we propose a way to extend the QRE model to A+B auctions. Let $\mathcal{P} = \{p_1, p_2, ..., p_{\#p}\}$ and $\mathcal{Q} = \{q_1, q_2, ..., q_{\#q}\}$ represent the set of all possible values of price and quality bidders can choose, respectively. The set of all possible values of private information z is $\mathcal{Z} = \{z_1, z_2, ..., z_Z\}$. A bidding strategy summarized by $\mathbf{S}(p, q|z)$ measures the probability that a bidder will choose any combination of (p, q) bid conditional on his private information z. For the strategy to be a well-defined probability measure, no price and quality bid should receive a negative probability, and conditional on any private information z, the sum of the probability of all possible price and quality combinations being chosen equals to one.

$$\mathbf{S}(p, q|z) \geq 0 \text{ for all } p \in \mathcal{P}, \, q \in \mathcal{Q}, \text{ and } z \in \mathcal{Z}, \quad (2.12)$$

and

$$\sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mathbf{S}(p, q|z) = 1 \text{ for any } z \in \mathcal{Z}. \quad (2.13)$$

If all bidders follow the same bidding strategy $\mathcal{B}(p, q|z)$, the probability that bidder j wins the auction with a bid (p, q) satisfies

$$x(s) = \left[\sum_{z \in \mathcal{Z}} \sum_{s' < s} \mathbf{S}(p, q|z) f(z) \right]^{N-1}, \text{ where } s = aq - p. \quad (2.14)$$

In the QRE model, the expected utility of bidder j with private information z_j is,

$$u(p, q; z_j) = (p - F - \frac{q^2}{2z_j})^\theta \times x(s) + \varepsilon = \pi(p, q|z_j)^\theta x(s) + \varepsilon, \quad (2.15)$$

where $s = aq - p$, and the optimization error ε is assumed to be i.i.d. with a CDF of $F(\varepsilon) = \exp[-\exp(-\lambda\varepsilon)]$, with a mean of γ/λ and a variance of $\pi^2/6\lambda^2$. The probability that bidder j chooses (p, q) conditional on a private value draw z_j in a A+B auction with

$N-1$ bidders follows

$$\sigma(p, q; z_j) = \frac{\exp\left(\lambda\pi(p, q|z_j)^\theta x(s)\right)}{\sum_{p'\in\mathcal{P}}\sum_{q'\in\mathcal{Q}}\exp\left(\lambda\pi(p', q'|z_j)^\theta x(s)\right)}. \tag{2.16}$$

According to McKelvey and Palfrey (1995), equilibrium exists, and the strategy $\mathbf{S}(p, q|z_j) = \sigma(p, q; z_j)$. As $\lambda \to 0$, the variance of the optimization error goes to infinity, meaning that when optimization is too costly, bidders would choose randomly. On the other hand, as $\lambda \to \infty$, the variance of the error term goes to zero, and the equilibrium converges to a Bayes-Nash equilibrium.

2.3 Experimental Design

For our experiments, we consider four treatments, varying the number of bidders, N, and the number of bidding dimensions. These are summarized in Table 2.1:

Table 2.1: Description of Treatments and Sample Size

Treatments	A+B auction	Price only
N=4 (1 human vs 3 automatans)	T1 (21 participants)	T2 (22 participants)
N=2 (1 human vs 1 automaton)	T3 (29 participants)	T4 (24 participants)

The streams of random production efficiency types were the same for all treatments (although the efficiency types were re-drawn after every period, every auction had the same N draws within a given round). Half of the participants were assigned to each treatment, and we used a between-subjects design.

Participants were matched in groups of N bidders for 50 rounds against N–1 automated bidders pre-programmed to play the Symmetric Bayes Nash Equilibrium strategy bids. Subjects were informed that "competitors are pre-programmed to place bids with the goal to obtain the highest profit possible." The experimental design aims to avoid social considerations (e.g., level-k theory) in the bidding process so that the participants were motivated to think and submit the best bids possible. The bidders were not allowed to

participate in more than one session nor more than one treatment. They were compensated based on their performance in terms of profits.

Experimental Parameterization. In our experiments, we considered the efficiency type z_j to be uniformly distributed between $L = 10$ and $U = 90$, elasticity of supply for quality η to be equal to unity, and the baseline cost F to be equal to 500. We fixed the buyer's quality preference factor to $a = 10$. These parameters provided bidders a numerical setting simple enough for them to handle, yet one that still captured the real-world features of the problem.

Because the computers were pre-programmed to place risk neutral SBNE bids, computers' strategy and distribution of bids given the parameterization follow:

$$\nu = \frac{1}{2}a^2 z - F = 50z - 500 \sim unif[0, 4000] \tag{2.17}$$

$$q^* = az = 10z \sim unif[100, 900] \tag{2.18}$$

$$s^* = b(\nu_j) = \frac{N-1}{N}\nu_j = \begin{cases} 37.5z - 375 \sim unif[0, 3000] & \text{if } N = 4 \\ 25z - 250 \sim unif[0, 2000] & \text{if } N = 2 \end{cases} \tag{2.19}$$

$$p^* = aq^* - s^* = \begin{cases} 62.5z + 375 \sim [1000, 6000] & \text{if } N = 4 \\ 75z - 250 \sim [500, 6500] & \text{if } N = 2 \end{cases} \tag{2.20}$$

In terms of scores, optimal bidding suggests that the SBNE scoring bid for the TS case with these parameters should be $(N-1)/N = 75\%$ of the bidder's pseudo-type ν_j (optimal bid-shading) in the case of 4 bidders and $(N-1)/N = 5\%$ in the corresponding case of 2 bidders.

In Treatments T1 and T3, bidders had to submit a simultaneous bid on price p and quality q based on their observed private efficiency type z. In Treatments T2 and T4, bidders observed the efficiency and the quality level chosen on their behalf based on their

efficiency ($q_j^* = 10z_j$). Bidders only need to submit a bid on price p. This design effectively reduced the dimensionality of the problem to a price-only procurement auction.

Procedure. A total of 96 subjects participated in the 4 treatments, and each session involved only one treatment. Each subject was compensated a proportion of their total earnings for all 50 rounds plus a \$10 participation fee. The average pay was approximately \$23, and each subject participated only in one session. Payment was made in private at the end of the session, and cash was the only incentive offered. Quality units, revenues, costs and profits were expressed in experimental currency units (ECUs). All ECUs earned throughout the session were transformed into U.S. dollars at a pre-specified rate of 1,100 ECUs per \$1.

All sessions were conducted at a major public research university, using the subject pool associated with the business school. Subjects privately read the printed instructions (see Appendix A for more details). Then, the researcher read the same instructions aloud and showed screen captures as examples. Approximately 10 minutes were used for reading instructions and answering questions, and 80 minutes were used for the decision task.

We programmed the experimental interface using the zTree system (Fischbacher 2007). The bidding screen included the realization of efficiency type z_j as well as the parameters for fixed cost $F = 500$ and weight $a = 10$. It also included a table showing all relevant information from past periods (own offered price and quality, winning price and quality, profits, etc.). The cost function and scoring function equations were provided on the instructions and the bidding screen. Because the decision task is computationally challenging, we also provided the participants with a decision support tool that automatically calculated the appropriate costs and earnings for winning the auction using the efficiency type and their inputs of price and quality. Because of this tool, the bidders were prevented from incurring computational mistakes that accidentally would result in bids generating a negative profit. After subjects in an auction submitted their offers, they saw a results screen showing if they had won, their own offered price and quality, the auction's winning price and quality, their resulting profits, and a table showing the history of all their rounds.

Subjects were made aware they were matched with automated computers in every round.

2.4   Reduced Form Analysis

In this section, we treat bidders' private information about their cost structure and the derived SBNE strategy as given, and we study if the models discussed in section 2.2 are able to describe their behavior. We first drop the data from the first 10 rounds of auctions to exclude the effect of learning on our analysis. Table 2.2 shows the summary statistics for the bidders' behavior over the final 40 auctions for each of the treatments (a more detailed summary of each bidder's behavior can be found in Appendix B). Across all four treatments, we can see that an average bidder's percentage of winning is higher than the risk-neutral SBNE predicted probability of winning, while his average requested profit per auction (not conditional on winning) is smaller than the profit that can be realized by following the risk-neutral SBNE strategy.[1] This result seems to support the risk-averse theory. Comparing Treatment 1 with Treatment 2 and Treatment 3 with Treatment 4, we can see that when the auction is in A+B format, both the percentage of winning and the requested profit is smaller than when bidders need to bid in only one dimension (quality is optimally given). This result provides evidence that when the complexity of the auction format increases, bidders are more likely to make errors in determining their optimal strategy. The increased difficulty in decision making in the A+B format is also seen in the score bids. However, the average score bids in the A+B auctions are smaller than those predicted by the risk-neutral SBNE strategy, contrasting with the prediction of risk aversion. Before we analyze the bidders' behaviors more closely, we examine the extent to which complexity in the A+B design affects its ability to efficiently allocate resources.

Unlike Table 2.2 that shows the bidder's behavior over all auction runs regardless of winning, Table 2.3 evaluates the performance of different designs by examining how human

---

[1]The RN SBNE column shows the auction outcome if a bidder follows the risk-neutral SBNE strategy based on the private information draws in the experiment. The numbers are slightly different from the theoretical mean because of the limited number of random draws from the distribution. For example, the theoretical probability of winning for 2 bidders is 50%; however, in this experiment, it is 45%.

Table 2.2: Summary Statistics for Bidding Behavior per Treatment

| | N=4 | | | | | N=2 | | | | |
| | T1 | | T2 | | | T3 | | T4 | | |
	Mean	SD	Mean	SD	RN SBNE	Mean	SD	Mean	SD	RN SBNE
% wins	32.5%	8.3%	37.0%	5.8%	25.0%	47.1%	12.7%	54.6%	8.6%	45.0%
avg price	4741.28	703.57	3582.15	196.27	3629.77	4670.62	761.68	4038.99	176.22	4155.72
avg quality	589.51	47.74	520.76	0	520.76	562.99	60.96	520.76	0	520.76
avg score	1153.8	444.94	1625.47	196.27	1577.86	959.32	391.27	1168.63	176.22	1051.91
avg profit	163.21	57.71	200.71	48.61	206.05	419.12	144.07	568.33	84.53	662.82
# of bidders	21		22			29		24		

Table 2.3: Evaluation of Auction Design Competition and Complexity on Human Winning

		T1	T2	T3	T4
	Buyer Surplus	2507.63	2480.35	1756.66	1646.32
Actual	Bidder Profit	502.19	541.79	890.44	1041.22
	Total Surplus	3009.82	3022.14	2647.10	2687.54
	Buyer Surplus	2339.25	2300.32	1447.54	1409.80
Optimal	Bidder Profit	779.75	766.77	1447.54	1409.80
	Total Surplus	3119.00	3067.09	2895.07	2819.59

decision making errors impact auction outcome. This table excludes data from the auction runs where the human bidder loses.[2] By comparing the total surplus between the actual and the optimal outcomes across treatment, A+B auctions appear to perform well in choosing the most efficient supplier. However, the actual bidder profit is lower, and the buyer surplus is higher than optimal comparing A+B auction with its price only counterpart. However, our evaluation based on this experiment is limited because in the field, quality is never optimally given based on bidder's private information. A better comparison is between the A+B auction and the price only auction with a minimum quality threshold. Nevertheless, our result suggests that the A+B auction is efficient despite increased bidding complexity.

2.4.1 Symmetric Bayesian Nash Equilibrium

In this section, we examine whether bidder behavior follows a symmetric Bayesian Nash Equilibrium strategy. Human bidders face bids submitted by computers following the

---

[2]Buyer surplus is measured by score bid; total surplus is the sum of buyer surplus and bidder profit.

SBNE strategy. Assume that human bidders have a rational expectation of the distribution of other bidders' (computers) score bids, then $G(s)$ is the CDF of an uniform distribution with support $[0, 3000]$ for Treatment 1 and 2 and $[0, 2000]$ for Treatment 3 and 4. Therefore, based on Equation 2.9 and 2.11, the risk-neutral SBNE strategy for score bid for human bidders is

$$s_j^* = \frac{N-1}{N}\nu_j, \qquad (2.21)$$

and the risk-averse SBNE strategy for score bid is

$$s_j^* = \frac{N-1}{N-1+\theta_j}\nu_j. \qquad (2.22)$$

Given this function form, we estimate θ_j using the following regression,

$$s_{jt} = \alpha_j + \beta_j\nu_{jt} + \varepsilon_{jt}. \qquad (2.23)$$

Then we are able to calculate the risk-aversion rate $\theta_j = \frac{1-\hat{\beta}_j}{\hat{\beta}_j}(N-1)$. We run Equation 2.23 separately for all bidders across treatment. Table 2.4 reports the mean and standard deviation for each bidder's intercept, slope, and the implied risk aversion. First, we see that the average α_j across all treatments is negative, meaning that when pseudo-type ν is small, human bidders tend to submit negative score bids. This result contradicts the SBNE theory because the score distribution has non-negative support, and a negative score bid is strictly dominated. In addition, we see the estimated average risk aversion varies greatly across treatments, ranging from -0.50 to 0.79. It is unlikely that a different auction format changes the risk preferences of bidders. To conclude, we do not find evidence from our data that the SBNE strategy predicts bidder behaviors accurately.

Table 2.4: Score Bid Regression Results

	T1		T2		T3		T4	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Intercept	-1542.40	1007.26	-280.84	604.54	-731.1	685.37	-96.34	409.97
slope	1.28	0.32	0.91	0.22	0.8	0.28	0.60	0.17
risk aversion	-0.50	0.72	0.46	0.65	0.4	0.49	0.79	0.54

2.4.2 Quantal Response Equilibrium

Since the scores human bidders face are uniform-distributed, based on Equation 2.16, in Treatment 1 and 3, human bidders' QRE strategy follows

$$\mathbf{S}(p, q; z_j) = \frac{\exp\left(\lambda \pi(p, q|z_j)^\theta x(s)\right)}{\sum_{p' \in \mathcal{P}} \sum_{q' \in \mathcal{Q}} \exp\left(\lambda \pi(p', q'|z_j)^\theta x(s')\right)}, \tag{2.24}$$

where $s = aq - p$, and in Treatment 2 and 4

$$\mathbf{S}(p; z_j) = \frac{\exp\left(\lambda \pi(p|z_j)^\theta x(s)\right)}{\sum_{p' \in \mathcal{P}} \exp\left(\lambda \pi(p'|z_j)^\theta x(s')\right)}, \tag{2.25}$$

where $s = a(az) - p$.

We can see from Equation 2.24 and 2.25 that the probability each possible bid ((p, q) in A+B auction and p in price only auction) chosen by an individual depends on (i) the expected utility that the bid generates and (ii) the number of possible bids. If a bid generates low expected utility, then it is likely to be chosen. On the other hand, holding the expected utility generated by a bid constant, if other more possible bids are available, it reduces the probability of this bid being chosen. Therefore, the increased complexity in A+B auction is directly reflected in the enlarged dimension of possible bids. In the price only auction (Treatment 2 and 4), the number of possible bids is only $\#p$, while in A+B auction (Treatment 1 and 3), the number of possible bids is $\#p \times \#q$. This prediction is reflected in Figure 2.1, where we compare the relationship between the actual score bids and the private information z across all 4 treatments. First, we can see that the actual bids do

55

not follow the SBNE strategy perfectly; rather they are dispersed around it, confirming our conclusion in the previous subsection. Second, we see that across all 4 treatments, when the efficiency draws are lower, the score bids are more noisy (wider confidence interval), and when efficiency draws are higher, the observed score bids are less noisy and follow the SBNE strategy closer. In addition, comparing Treatments 1 and 3 with Treatments 2 and 4 indicates that the deviation in the score bids from the SBNE strategy is larger for the A+B auction than for the price only auction, a direct result from the higher dimension of the bids. Furthermore, the discrepancies in score bids seem to be larger in Treatments 1 and 2 than in Treatments 3 and 4 because the higher number of bidders results in lower expected profit and, thus, more noise in decision making.

Figure 2.1: Observed Score Bids and Risk-Neutral SBNE Score Bids
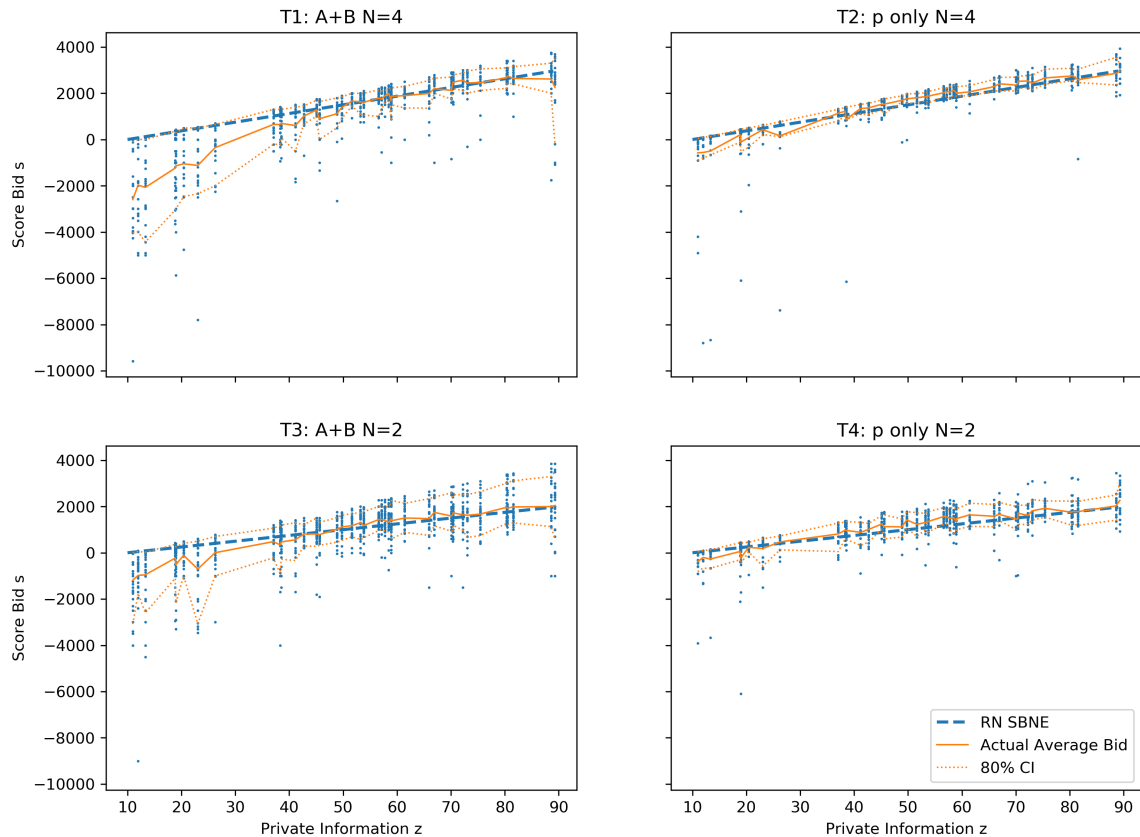
Another impotent feature shown in Figure 2.1 is that the relationship between the

actual score bids and the private information is non-linear. In all 4 treatments, when efficiency levels are lower, we observe more underbidding in the scores. This feature has not been observed by other experimental papers on first price sealed-bid auction. A potential reason for this result is the difference in experimental design. In the other experimental papers, the bidders are assigned a private value, and they bid the price they want to pay. Because price cannot be negative, the possible price bid are bounded between the bidder's private value and zero. However, according to the instruction for our experiment (Appendix A), bidders can choose any price value between 0 to 10,000 and a quality value greater than 0. In addition, the software automatically calculates and displays the bidders' costs and earnings based on their choice of price and quality bid. It also addresses that if the price is below cost, a bidder will lose money. Consequently, bidders will not choose a combination of price and quality that generates a negative profit. However, since the resulting score from a price and quality bid is not calculated and displayed by the system, bidders are likely to choose a negative score. Although negative scores are dominant in the SBNE, in the QRE, because of the optimization error, the probability of bidders choosing a negative score is non-negative. Figure 2.2 shows the expected utility when the efficiency draw is low ($z = 20$) and the highest expected utility is only 0.24 ECU. The area below the lower boundary of the 0 contour shows the price and quality combination that yields negative profits. Because the expected utility is close to 0, bidders only choose almost uniformly from the area lies in the top-left of the lower 0 contour boundary, suggesting we see underbidding when the efficiency level is low.

Figure 2.3 plots the expected utility when the efficiency level $z = 50$. The price and quality combination that generates the highest expected utility is 3500 and 500. When we hold the price bid constant and change the quality bid, the deviation in the latter from the optimal value results in a larger utility loss than when we fix the quality bid and adjust the price bid. As a result, under QRE, we expect to see a larger deviation from the optimal in the price bid than in the quality bid. As shown in Figure 2.4, the price bids of the human bidders are very noisy, more so in Treatment 1 than in Treatment 3, where the

Figure 2.2: Price and Quality Bids Contour $z = 20, N = 4$
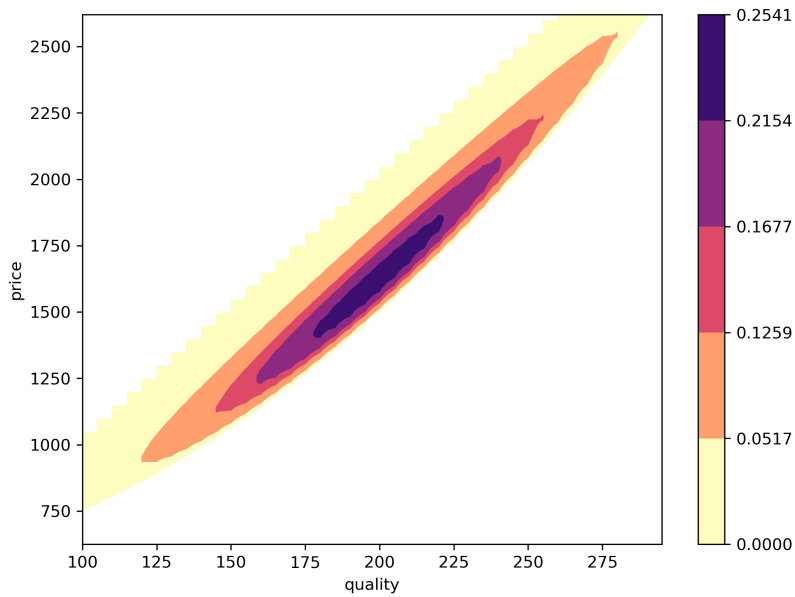
Figure 2.3: Price and Quality Bids Contour $z = 50, N = 4$
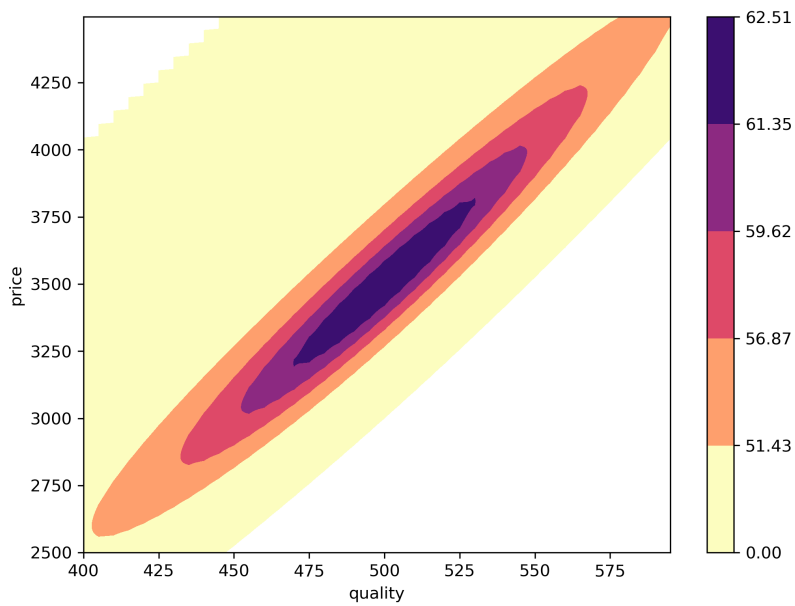
larger number of bidders reduces the incentive to choose optimally. In contrast, as shown in Figure 2.5 and Table 2.5, bidders do reasonably well at choosing the quality level although they tend to choose higher than optimal quality bids. There is no significant difference in

quality bids in Treatments 1 and 3, as the probability of winning varies by competition intensity. This result is consistent with theory.

Figure 2.4: Observed Price Bids and SBNE Price Bids
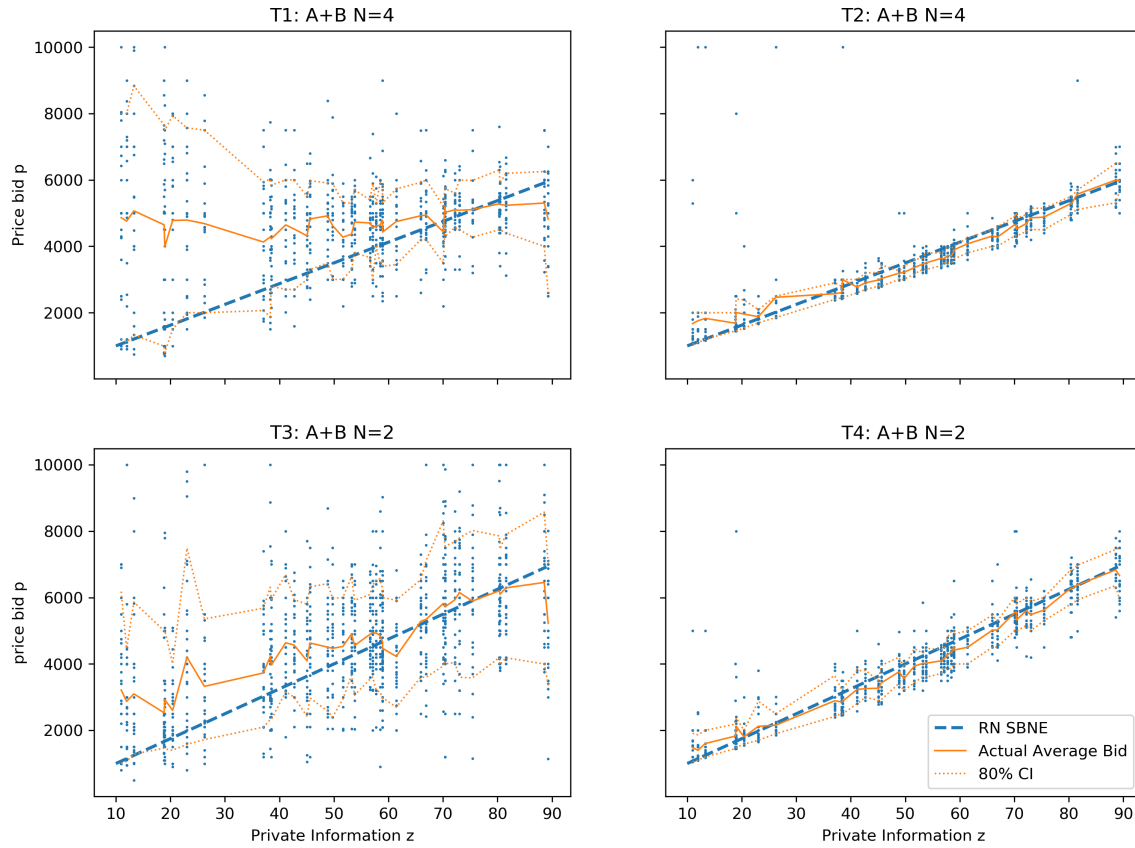
Table 2.5: Quality Bid Regression

	T1 I	T1 II	T3 I	T3 II
Intercept	217.9236***		116.2660***	
	(9.8317)		(10.0058)	
z4	7.1354***	10.7043***	8.5783***	10.4824***
	(0.1744)	(0.0842)	(0.1774)	(0.0719)
N	840	840	1160	1160
R2	0.67	0.95	0.67	0.95

Figure 2.5: Observed Quality Bids and SBNE Quality Bids
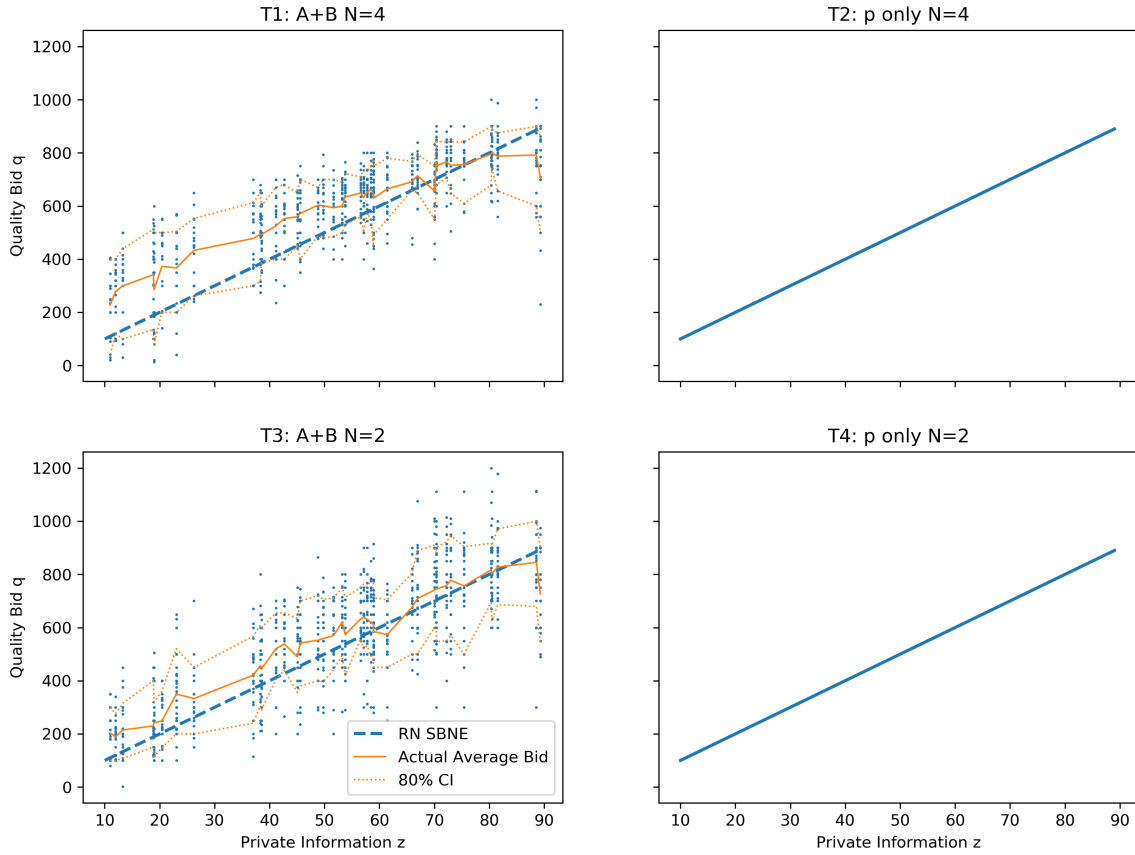
2.5 Structural Analysis

In this section, we treat bidders' private information about their cost structure as unknown. We apply structural estimation methods based on the models discussed above to estimate bidders' private information. We further evaluate the prediction accuracy of different models.

2.5.1 A Probability Weighting/Prospect Theoretical Framework Under Risk Neutrality

Bidders decide among risky alternatives. If we assume that human bidders are risk neutral but have difficulty assessing the probability associated with the different events linked to their decisions, a structural estimate of the probability measure of the pseudo-

types implicit would yield a measure of this probability.

Since structural nonparametric estimates of the pseudo-type and its distribution are identified, as shown by Quiroga (2015)'s extension of the Guerre et al. (2000)'s identification results for multiattribute bids, we can use the resulting estimates as behavioral measures of both the implicit valuation and the associated probability weights. As a result and given that the losses realm is precluded by our experimental design, all that Prospect Theory (Tversky and Kahneman 1992) can prescribe accurately is the distortion in the probability weights, and if this is the case, the Kernel-based nonparametric estimates of the PDF and CDF of the pseudo-types would, in theory, be the most appropriate measures available for this phenomenon.

First stage estimates: Empirical Pseudo-types. As a first step in the consecution of the distributions, we obtain nonparametric estimates of the empirical pseudo-types from the following relationship:

$$\hat{\nu}_j = s_j + \frac{1}{N-1}\frac{\hat{G}(s_j)}{\hat{g}(s_j)} \tag{2.26}$$

While this result was derived under a fully rational risk neutral framework, it can be extended to the perceived probability $\hat{G}(s_j)$ used by subjects in their choice evaluations, itself a behaviorally distorted measure of the true probability measure for the actual pseudo-types, $F_\nu(\nu)$, which will be obtained in the second stage estimates.

Figure 2.6 displays plots of the induced values (pseudo-types) against the kernel estimated values (empirical pseudo-types) for each treatment. We used the standard Guerre et al. (2000) sample trimming procedure to address the boundary effects inherent in kernel density estimation. More specifically, we discarded the value estimates corresponding to the bids that lay in $[b_{min}, b_{min} + h] \cup [b_{max} - h, b_{max}]$, where h is the bandwidth determined by the Hardin rule of thumb. Throughout this paper, we use sample trimming as a default method for kernel estimation. In addition to boundary trimming, we also discarded the top and bottom 5% of the value estimates in order to obtain results in the strict interior of the

distributions.

These graphs show how heterogeneous the implicit (i.e., estimated) valuation is across bids, especially when compared to the induced valuation (which is denoted by the solid black line). T2 (price only, 3 competitors) is once again the most stable around the induced pseudo-types line.

Figure 2.6: Plots of Induced Pseudo-types and Estimated Valuations
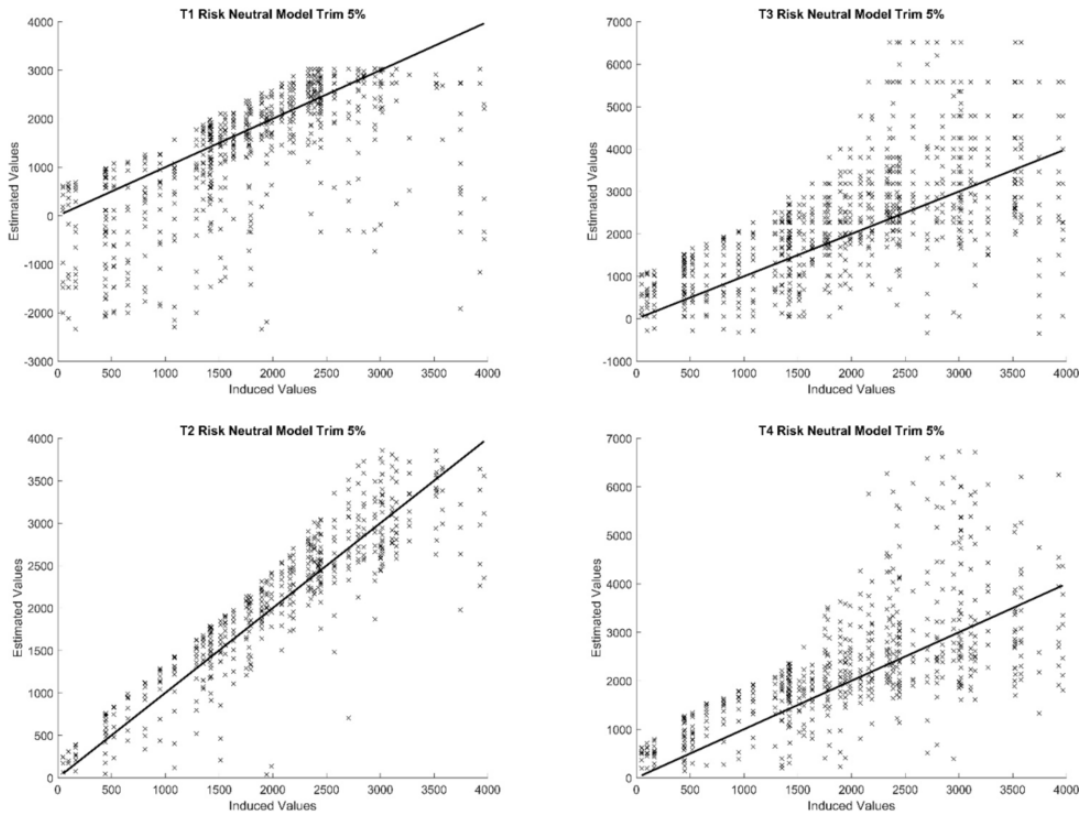
Second stage estimates: Empirical Probabilities. With the first stage estimates, we can again use kernel methods to retrieve the probability densities and cumulative distribution functions implicit in the suppliers' score bid decisions.

Figure 2.7 shows the estimated probability weighting functions in terms of CDFs. The CDFs are evaluated between the 5th and 95th percentiles of the induced pseudo-types. Confidence intervals were calculated using bootstrap, conducted by randomly choosing 15

bidders from all bidders in a treatment 1,000 times. The CI of estimated values is con-
structed using the subsamples at the 5th and 95th percentile. CDFs of the estimated values
are calculated using the average of the estimated values of all subsamples.

Figure 2.7: Probability Weighting Functions Estimated via Nonparametric CDFs of Em-
pirical Valuations
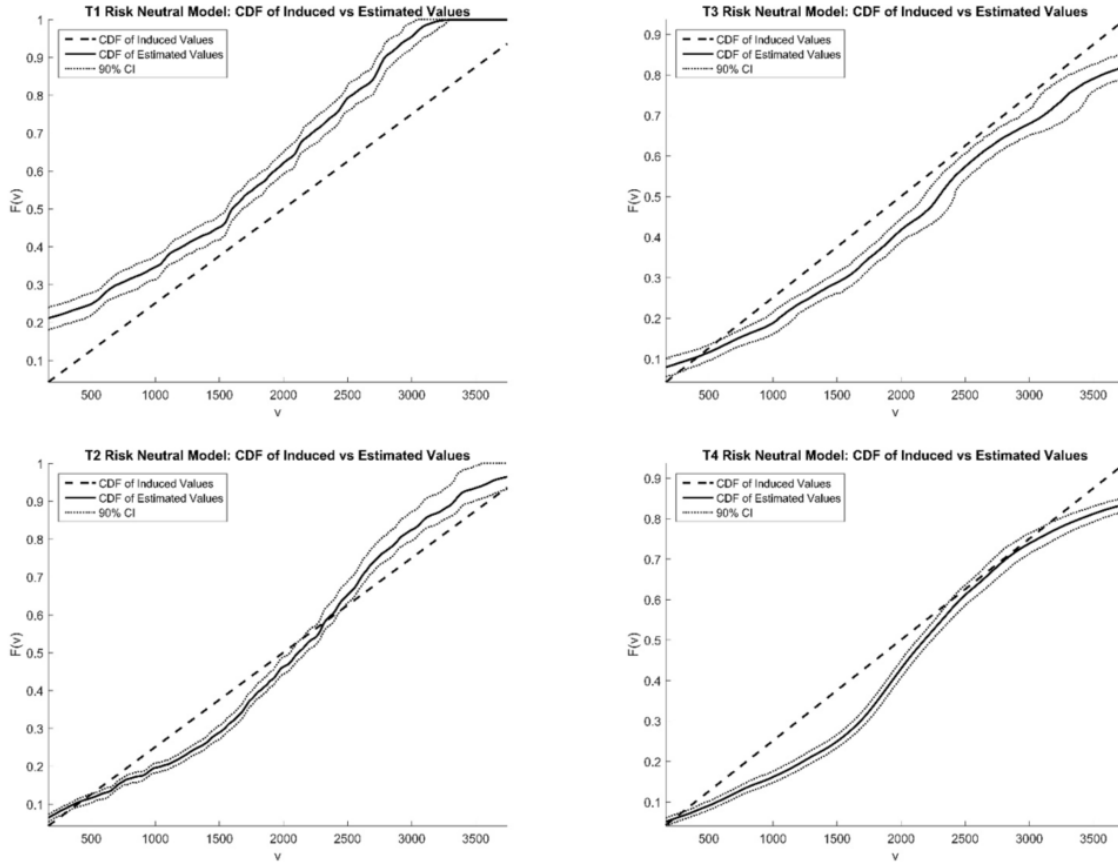
In T1, the CDF of the estimated values is higher than the CDF of the induced
pseudo-types, implying first-order stochastic dominance of the pseudo-types over the esti-
mated values. This is a direct effect of the combination of higher complexity (bids in two
dimensions) and high competition. The results from T3, except for the case with very low
values, are opposite from the first-order stochastic dominance, which from parallelism is
inferred as a competition effect.

For T2 and T4, no clear dominance results are implied, and the shape of the curve

underweights probabilities for lower values with respect to the induced probability, and overweights probabilities for higher values, results that contradict what Tversky and Kahneman (1992) suggested would be the traditional psychological bias in probabilities (which is predicted to have an inverse S-shape as opposed to the S-shape observed in treatments T2 and T4). Our results for these two treatments, hence, suggest increased sensitivity as opposed to the diminishing sensitivity hypothesis of Tversky and Kahneman (1992). It is not entirely unexpected, as earlier works (e.g. Schweitzer and Cachon 2000, Nagarajan and Shechter 2014) had already provided evidence of the problems that probability weighting bias/prospect theoretical frameworks face in rationalizing other decision problems under uncertainty such as the news vendor problem. Here, we cannot ignore that the probability measures used by bidders might be distorted and that this distortion may be the explanation for the bidding behavior far from the standard SBNE solution. However, we acknowledge that there is more work to be done in this respect.

Table 2.6 summarizes the estimated results. We evaluated goodness of fit of the nonparametric estimates of the Risk Neutral Model to the SBNE prediction using the Kolmogorov Smirnov (KS) statistic as used in Bajari and Hortacsu (2005),

$$KS = \sqrt{T} \max_{[\nu_0, \nu_1]} |\hat{F}(\nu) - F(\nu)|. \tag{2.27}$$

In addition, we also computed the norms L^1 and L^2 as measures of distance:

$$L^1 = \frac{1}{T} \sum_{j,t} |\hat{\nu}_{jt} - \nu_{jt}| \tag{2.28}$$

$$L^2 = \Big[\frac{1}{T} \sum_{j,t} (\hat{\nu}_{jt} - \nu_{jt})^2\Big]^{1/2} \tag{2.29}$$

64

Table 2.6: Estimation Results for the Nonparametric Kernel Risk Neutral Model

	T1	T2	T3	T4
KS Statistic (5th-95th percentile of value support)	0.20284	0.05163	0.18103	0.12653
p value	0.00000	0.10330	0.00000	0.00000
Reject Equality	Y	N	Y	Y
KS Statistic (25th-75th percentile of value support)	0.17588	0.05163	0.10137	0.08185
p value	0.00000	0.10330	0.00001	0.00059
Reject Equality	Y	N	Y	Y
Critical Value	0.04197	0.04101	0.03926	0.03926
L^1	836.69206	355.96426	1218.41531	1055.75457
L^1/m (value)	0.47398	0.17785	0.60046	0.51723
L^2	1425.38778	652.06852	2167.30129	2054.79150
L^2/m (value)	0.80748	0.32580	1.06810	1.00668

2.6 Conclusion

In this paper, we conducted an experiment of A+B auction and examined bidders' behavior. We found that bidder behaviors systematically deviate from the Bayesian-Nash equilibrium strategy. However, despite the increased noise in bidders' decision making due to increased complexity in the scoring rule, A+B auction performs well in terms of discovering the most efficient supplier. We further proposed a way to model complexity, by extending the Quantal Response Equilibrium to A+B auctions. We found that the systematic deviations of bidder behaviors from the Bayes-Nash prediction across different levels of competition and complexity in the auction format are consistent with the predictions of the Quantal Response Equilibrium model. We also evaluated the ability of the Bayesian-Nash Equilibrium model to estimate bidders' private values using a nonparametric method, but the measures performed poorly as expected. We are hopeful that structural methods based on the Quantal Response Equilibrium model will generate reasonable estimates.

Appendices

Appendix A   Experiment Instruction

Instruction

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you will earn money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will come to your station and answer it. We ask that you not talk with one another for the duration of the experiment.

Government agencies or companies often ask suppliers to participate in a Procurement Auction by submitting a bid to supply equipment or construct a building, a bridge or a highway. These bids are often at a combination of price, quality and possibly other factors.

In this experiment, you will be in the role of a **Supplier (Bidder)** who participates in 40 Procurement Auctions, competing against **three (3) computerized** competitors (there are 4 suppliers in total). Your competitors are pre-programmed to place bids with the goal to obtain the highest profit possible.

The session consists of **50** separate auction rounds, and you will submit a bid in each round. The unit of exchange in all the transactions is called experimental currency units (ECUs). At the end of the session, your earnings in ECUs will be converted to US dollars at a pre-specified rate, and paid to you in cash.

How you earn money

Each bid you make consists of (A) a money price you request in ECUs ("Your Price"), and (B) a quality level ("Your Chosen Quality"). The buyer has established a scoring rule to compare bids submitted by each supplier. Supplier Score increases with Chosen Quality, and decreases with Price:

Supplier Score = (Buyer Quality Weighting * Your Chosen Quality) – Your Price

"Buyer Quality Weighting" represents the evaluation the buyer makes about the quality level offered by any supplier. It is a positive constant, known and common to all

bidders, and equal to 10.

The highest score among all four competitors wins the contract (Any ties in supplier scores are broken at random). You earn money each time you win a contract at a good price and appropriate quality. If you win the contract:

$$\text{Your Earnings} = \text{Your Price} - \text{Your Cost}$$

If you don't win the contract, your profit, your earnings and your cost will be 0 regardless of your bid.

Only if you win the contract, you incur a production cost ("Your Cost") to produce the good at Your Chosen Quality. This cost varies based on Your Efficiency and Fixed Costs:

$$\text{Your Cost } = \left(0.5 \times \frac{(\text{Your Chosen Quality})^2}{\text{Your Efficiency}}\right) + \text{ Fixed Costs}$$

You will see Your Efficiency before deciding how much to bid, but not the efficiency of any other competitor. For each bidder, the higher the efficiency, the cheaper it is to offer the chosen quality level. Your Efficiency will change from round to round; it is independent from your competitors' efficiencies, and is unrelated to your efficiency in any other auctions. Your Efficiency is a random number within the range $[10, 90]$, with all values in that range being equally likely. Fixed Costs are the same and are known for all bidders before each auction, but are only incurred by the winner.

For each auction, enter a number for Your Price and another for Your Chosen Quality in the boxes on your computer screen, and then click the "Make Offer" button. For price decisions you can chose any number between 0 and 10,000. For quality levels, you can choose any number from 0 and above. Your decisions can have up to two decimals.

With the above information, you will be able to calculate Your Cost and Your Earnings in each round in case you win. Our software has an embedded simulator you can use for this. Each bidder knows only their own cost, but not the cost of any other bidder.

Note: If a supplier bids his/her Price below his/her Cost, and wins the auction, he/she loses money. Therefore, carefully choose your price and quality. Hint: Use the Simulator tool for this purpose.

Information you will see at the end of each auction

At the end of each auction you will see the following information:

- Your own efficiency and bid (price and quality) in this auction.

- Your Cost (which will only be different from 0 if you win).

- The winning bid (price and quality).

- Your earnings from the auction.

You will also have access to this information for all past auctions.

How you will be paid

At the end of the session you will see the final screen summarizing your earnings for the session. This screen will calculate your net profits from the 50 auctions, convert them to US dollars at the rate of 1,100 ECU per $1, and add them to your $5 participation fee.

Please use this information to fill out your check-out form and wait quietly until the experimenter calls you to come to the front of the room and be paid your earnings in private and in cash. After you have been paid, you may leave the laboratory.

Appendix B   Summary Statistics on Individual Bidder's Behavior per Treatment

Table 7: Treatment 1 Summary Statistics on Individual Behavior

	treatment	% wins	avg price	avg quality	avg score	avg profit
10101	1	42.5%	4531.25	609.98	1568.50	161.21
10102	1	30.0%	5000.21	615.16	1151.37	38.85
10201	1	32.5%	4916.52	621.62	1299.72	186.54
10202	1	40.0%	4328.08	585.14	1523.29	209.58
10301	1	25.0%	4978.93	579.26	813.63	151.23
10302	1	17.5%	5459.55	596.68	507.20	120.53
10401	1	22.5%	3742.90	521.83	1475.38	191.43
10402	1	20.0%	5784.25	612.45	340.25	126.58
10501	1	42.5%	4927.12	621.85	1291.38	95.57
10502	1	42.5%	5120.75	633.42	1213.50	253.57
10503	1	40.0%	3836.18	553.62	1700.08	208.91
10601	1	40.0%	5596.08	682.50	1228.92	37.42
10602	1	37.5%	4007.25	563.80	1630.75	163.69
10603	1	20.0%	5840.10	593.17	91.64	186.89
10701	1	35.0%	3863.89	549.55	1631.60	220.98
10702	1	30.0%	5422.50	636.00	937.50	211.63
10703	1	30.0%	4151.20	505.45	903.30	156.94
10801	1	25.0%	4662.88	595.25	1289.62	180.88
10802	1	32.5%	5391.85	623.43	842.41	241.48
10803	1	42.5%	4423.18	601.65	1593.33	146.22
10804	1	35.0%	3582.25	477.88	1196.50	137.33
Mean		32.5%	4741.28	589.51	1153.80	163.21
SD		8.3%	703.57	47.74	444.94	57.71

Table 8: Treatment 2 Summary Statistics on Individual Behavior

	treatment	% wins	avg price	avg quality	avg score	avg profit
20101	2	45.0%	3391.00	520.76	1816.62	112.90
20201	2	32.5%	3596.42	520.76	1611.20	234.79
20202	2	30.0%	3884.25	520.76	1323.37	106.40
20301	2	40.0%	3819.38	520.76	1388.25	189.89
20302	2	32.5%	3617.50	520.76	1590.12	208.89
20401	2	25.0%	3786.30	520.76	1421.32	209.04
20402	2	37.5%	3691.25	520.76	1516.37	205.70
20403	2	35.0%	3554.38	520.76	1653.25	246.91
20404	2	40.0%	3494.88	520.76	1712.75	121.02
20405	2	32.5%	3424.56	520.76	1783.06	102.51
20501	2	32.5%	3641.68	520.76	1565.95	258.87
20502	2	42.5%	3450.00	520.76	1757.62	249.16
20503	2	27.5%	3730.48	520.76	1477.15	223.20
20504	2	32.5%	4143.23	520.76	1064.40	196.66
20505	2	42.5%	3442.00	520.76	1765.62	236.54
20601	2	42.5%	3388.05	520.76	1819.57	179.25
20602	2	37.5%	3523.00	520.76	1684.62	217.79
20603	2	45.0%	3328.50	520.76	1879.12	183.40
20604	2	42.5%	3423.30	520.76	1784.32	219.57
20605	2	42.5%	3445.00	520.76	1762.62	222.29
20701	2	40.0%	3537.25	520.76	1670.37	249.74
20702	2	37.5%	3495.00	520.76	1712.62	241.04
Mean		37.0%	3582.15	520.76	1625.47	200.71
SD		5.8%	196.27	0.00	196.27	48.61

Table 9: Treatment 3 Summary Statistics on Individual Behavior

	treatment	% wins	avg price	avg quality	avg score	avg profit
30101	3	45.0%	5758.75	636.05	601.75	335.92
30102	3	65.0%	4281.82	576.33	1481.42	55.53
30103	3	27.5%	4890.42	567.15	781.08	424.89
30104	3	70.0%	4154.38	590.45	1750.12	93.16
30105	3	32.5%	5301.42	556.20	260.56	348.54
30201	3	35.0%	6151.90	675.32	601.33	450.89
30202	3	32.5%	5176.93	549.67	319.74	353.14
30203	3	37.5%	3897.49	472.50	827.49	288.73
30301	3	40.0%	4482.00	534.33	861.25	560.25
30302	3	57.5%	4560.38	577.22	1211.88	532.27
30303	3	40.0%	3755.62	451.88	763.12	447.51
30304	3	67.5%	3774.88	529.12	1516.38	232.10
30305	3	50.0%	4153.50	530.62	1152.75	566.59
30401	3	37.5%	4755.25	565.38	898.50	516.52
30402	3	25.0%	5989.70	619.25	202.80	382.24
30403	3	40.0%	4451.48	514.83	696.78	587.69
30404	3	55.0%	5010.00	620.01	1190.12	339.97
30405	3	42.5%	6048.68	692.35	874.82	290.53
30406	3	50.0%	5303.75	643.50	1131.25	506.87
30501	3	40.0%	4502.58	533.85	835.92	414.77
30502	3	42.5%	4599.50	549.50	895.50	549.54
30503	3	57.5%	3980.88	529.62	1315.38	583.79
30504	3	42.5%	4290.00	486.25	572.50	397.41
30505	3	45.0%	5620.00	659.25	972.50	495.01
30601	3	47.5%	4711.25	566.50	953.75	575.14
30602	3	70.0%	3388.75	506.80	1679.25	233.08
30603	3	70.0%	3508.50	479.25	1284.00	519.70
30604	3	47.5%	4895.75	585.65	960.75	509.54
30605	3	52.5%	4052.55	528.00	1227.45	563.19
Mean		47.1%	4670.62	562.99	959.32	419.12
SD		12.7%	761.68	60.96	391.27	144.07

Table 10: Treatment 4 Summary Statistics on Individual Behavior

	treatment	% wins	avg price	avg quality	avg score	avg profit
40101	4	62.5%	3750.62	520.76	1457.00	481.37
40102	4	50.0%	3921.72	520.76	1285.90	275.70
40103	4	47.5%	4164.40	520.76	1043.22	547.06
40104	4	55.0%	3961.25	520.76	1246.37	569.29
40105	4	37.5%	4278.13	520.76	929.49	537.36
40201	4	57.5%	3941.38	520.76	1266.25	544.79
40202	4	65.0%	3994.35	520.76	1213.27	597.69
40203	4	52.5%	4223.50	520.76	984.12	595.75
40204	4	47.5%	3958.75	520.76	1248.87	514.49
40205	4	55.0%	4232.30	520.76	975.32	583.76
40301	4	62.5%	3982.08	520.76	1225.55	680.94
40302	4	60.0%	3933.65	520.76	1273.97	638.02
40303	4	47.5%	4174.89	520.76	1032.74	605.22
40304	4	72.5%	3724.80	520.76	1482.82	523.81
40401	4	57.5%	3991.48	520.76	1216.15	648.71
40402	4	65.0%	3753.75	520.76	1453.87	510.30
40403	4	42.5%	4211.48	520.76	996.15	547.37
40404	4	62.5%	3910.01	520.76	1297.61	564.57
40405	4	45.0%	4278.75	520.76	928.87	645.46
40501	4	55.0%	4206.05	520.76	1001.57	690.57
40502	4	65.0%	3885.00	520.76	1322.62	624.45
40503	4	47.5%	3991.25	520.76	1216.37	508.79
40504	4	50.0%	4252.50	520.76	955.12	643.27
40505	4	47.5%	4213.70	520.76	993.92	561.24
Mean		54.6%	4038.99	520.76	1168.63	568.33
SD		8.6%	176.22	0.00	176.22	84.53

Bibliography

Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics 38*(1), 119–126.

Asker, J. and E. Cantillon (2008). Properties of scoring auctions. *The RAND Journal of Economics 39*(1), 69–85.

Atkinson, B. (2008). On retail gasoline pricing websites: potential sample selection biases and their implications for empirical research. *Review of Industrial Organization 33*(2), 161–175.

Bajari, P. and A. Hortacsu (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy 113*(4), 703–741.

Branco, F. (1997). The design of multidimensional auctions. *The RAND Journal of Economics*, 63–81.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, 457–482.

Campo, S., E. Guerre, I. Perrigne, and Q. Vuong (2011). Semiparametric estimation of first-price auctions with risk-averse bidders. *The Review of Economic Studies 78*(1), 112–147.

Chandra, A. and M. Tappata (2011). Consumer search and dynamic price dispersion: an application to gasoline markets. *The RAND Journal of Economics 42*(4), 681–704.

Che, Y.-K. (1993). Design competition through multidimensional auctions. *The RAND Journal of Economics*, 668–680.

Cox, J. C., V. L. Smith, and J. M. Walker (1985). Experimental development of sealed-bid auction theory; calibrating controls for risk aversion. *The American Economic Review 75*(2), 160–165.

Cox, J. C., V. L. Smith, and J. M. Walker (1988). Theory and individual behavior of first-price auctions. *Journal of Risk and uncertainty 1*(1), 61–99.

De los Santos, B. (2018). Consumer search on the internet. *International Journal of Industrial Organization 58*, 66–105.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review 102*(6), 2955–2980.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2017). Search with learning for differentiated products: Evidence from e-commerce. *Journal of Business & Economic Statistics*, 1–16.

Dinerstein, M., L. Einav, J. Levin, and N. Sundaresan (2018). Consumer price search and platform design in internet commerce. *American Economic Review 108*(7), 1820–59.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory 104*(1), 247–272.

Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica 68*(3), 525–574.

Gupta, D., E. M. Snir, and Y. Chen (2015). Contractors' and agency decisions and policy implications in a+ b bidding. *Production and Operations Management 24*(1), 159–177.

Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics 37*(2), 257–275.

Honka, E. (2014). Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics 45*(4), 847–884.

Hortaçsu, A. and C. Syverson (2004). 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 119*(2), 403–456.

Kim, J. B., P. Albuquerque, and B. J. Bronnenberg (2010). Online demand under limited consumer search. *Marketing science 29*(6), 1001–1023.

Koulayev, S. (2013). Search with dirichlet priors: estimation and implications for consumer demand. *Journal of Business & Economic Statistics 31*(2), 226–239.

Koulayev, S. (2014). Search for differentiated products: identification and estimation. *The RAND Journal of Economics 45*(3), 553–575.

Lewis, G. and P. Bajari (2011). Procurement contracting with time incentives: Theory and evidence. *The Quarterly Journal of Economics 126*(3), 1173–1211.

Lewis, M. S. (2008). Price dispersion and competition with differentiated sellers. *The Journal of Industrial Economics 56*(3), 654–678.

Lewis, M. S. (2011). Asymmetric price adjustment and consumer search: An examination of the retail gasoline market. *Journal of Economics & Management Strategy 20*(2), 409–449.

Lewis, M. S. and H. P. Marvel (2011). When do consumers search? *The Journal of Industrial Economics 59*(3), 457–483.

Matsumoto, B. and F. Spence (2016). Price beliefs and experience: Do consumers beliefs converge to empirical distributions with repeated purchases? *Journal of Economic Behavior & Organization 126*, 243–254.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113–126.

McKelvey, R. D. and T. R. Palfrey (1995). Quantal response equilibria for normal form games. *Games and economic behavior 10*(1), 6–38.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2015). Consumer search and prices in the automobile market.

Nishida, M. and M. Remer (2018). The determinants and consequences of search cost heterogeneity: Evidence from local gasoline markets. *Journal of Marketing Research 55*(3), 305–320.

Quiroga, B. F. (2015). Nonparametric identification and estimation of score auctions in multiattribute procurement. *Pontificia Universidad Catolica de Chile, School of Management, working paper*.

Quiroga, B. F., B. B. Moritz, and D. R. Guide, Jr. (2019). Transparency in procurement: Revealed vs. concealed scoring rules in a+b sealed-bid auctions. *Pontificia Universidad Catolica de Chile, School of Management, working paper*.

Remer, M. (2015). An empirical investigation of the determinants of asymmetric pricing. *International Journal of Industrial Organization 42*, 46–56.

Riley, J. and W. F. Samuelson (1981). Optimal auctions. *American Economic Review 71*(3), 381–92.

Rothschild, M. (1974). Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy 82*(4), 689–711.

Stigler, G. J. (1961). The economics of information. *Journal of political economy 69*(3), 213–225.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty 5*(4), 297–323.

Wildenbeest, M. R. (2011). An empirical model of search with vertically differentiated products. *The RAND Journal of Economics 42*(4), 729–757.