Intermediation and Competition in Search Markets: An Empirical Case Study *

Tobias Salz[†]

July 31, 2020

Abstract

Intermediaries in decentralized markets can affect buyer welfare both directly, by reducing expenses for buyers with high search cost and indirectly, through a search-externality that affects the prices paid by buyers that do not use intermediaries. I investigate the magnitude of these effects in New York City's trade-waste market, where buyers can either search by themselves or through a waste broker. Combining elements from the empirical search and procurement-auction literatures, I construct and estimate a model for a decentralized market. Results from the model show that intermediaries improve welfare and benefit buyers in both the broker and the search markets.

Keywords: Search Cost, Intermediation, Decentralized Markets

JEL Classification Codes: L13, D43, D44, L97, L81

*I am indebted to my advisors John Asker, Guillaume Frechétte, Kei Kawai, and Alessandro Lizzeri for their guidance and support. I am extremely grateful to Isabelle Perrigne and Quang Vuong for helpful discussions. Luis Cabral, Michael Dickstein, Klaus-Peter Hellwig, Kate Ho, Jean-Francois Houde, Myrto Kalouptsidi, Yunmi Kong, John Lazarev, Konrad Menzel, Joseph Mullins, Ariel Pakes, Maryam Saeedi, Maher Said, Kevin Thom, Allan Collard-Wexler, Larry White, Sevgi Yuksel, Jidong Zhou, as well as seminar participants at Caltech, Carnegie Mellon University, Columbia, Cornell, Duke, Harvard, LSE, Microsoft Research Cambridge, MIT, Northwestern, the New York Fed, NYU, Princeton, Stanford, UCLA, USC, UPenn, UT Austin and Yale SOM, provided invaluable feedback. The research in this paper was conducted with data provided by New York City's Business Integrity Commission. The data has been modified to preserve the anonymity of all market participants. In addition, results have been reviewed to ensure that no confidential information is disclosed. The research results and conclusions expressed are those of the author and do not necessarily reflect the views of the Business Integrity Commission. I want to thank Thi-Mai Anh Nguyen for exceptional research assistance. All remaining errors are mine.

† Department of Economics, Massachusetts Institute of Technology, 50 Memorial Dr, Cambridge, MA 02142. tsalz@mit.edu.

1 Introduction

In decentralized markets, a transaction typically requires a costly search for the cheapest or best-suited seller. These search frictions can increase sellers' market power and allow them to raise prices. At the same time, search frictions create a role for intermediaries that search on behalf of buyers and thereby mitigate this effect. Intermediaries are commonly found in markets with search frictions, but there is little systematic empirical evidence on how they influence market outcomes.[1] This paper uses new a dataset from the New York City trade waste industry to study study how intermediation affects market efficiency. In order to assess the impact of intermediaries on competition and welfare, I develop and estimate a model that combines elements from the empirical search and auction literatures.

Buyers that use the services of intermediaries directly benefit from access to a better search technology. At the same time, intermediaries can also benefit the other buyers through a search externality. Intermediaries are used by buyers with higher search cost. This selection changes the composition of buyers in the search market, increasing the proportion of those who can afford more comparison shopping. Therefore, sellers have an incentive to quote lower prices relative to a scenario without intermediaries. Results from the model suggest that this externality is quantitatively important for welfare.[2]

It is difficult to obtain data on the operation of decentralized markets due to the very nature of these markets. However, the trade waste industry in New York pro-

[1] Examples beyond housing and real estate include textile (http://www.economist.com/news/business/21657375-story-indias-biggest-maker-towels-and-their-journey-cotton-field-big-box (last accessed on 06/14/20)), advertising brokers, freight brokers (http://www.dat.com/∼/media/files/dat/resources/dat_broker_start-up_guide_rnd2.pdf (last accessed on 08/29/15)), literary agents (http://aaronline.org/ (last accessed on 06/14/20)), energy brokers (https://en.wikipedia.org/wiki/Energy_broker (last accessed on 06/14/20)). Spulber (1996b) estimates that 25% of U.S. economic activity is intermediation.

[2] This is reminiscent of the effect of informed buyers in Salop and Stiglitz (1977), when sellers can't distinguish between informed and uninformed buyers. Sellers know that some buyers search a lot and, therefore, need to lower their prices.

vides a rare opportunity to gain insight into such a market: as a measure against the historical entrenchment of organized crime in this industry, the city has established a regulator that monitors carters and collects data about many operational aspects of their activity. This study uses an anonymized panel of the universe of bilateral agreements between 100 private carters (the supply side) and more than 100,000 businesses (the demand side).

Several observations about the market indicate that search frictions play an important role. Price dispersion in the market is large, even after accounting for observable contract characteristics. Both sides of the market mainly feature smaller players. In the majority of zip codes, buyers can choose from more than 20 different waste carters. Buyers have the option to procure their contract through a waste broker and a significant portion of buyers makes use of this option.

My model combines elements from the empirical search and auction literatures. Sellers draw a match-specific cost for each buyer. Buyers can either contact carters directly at some cost per inquiry or delegate search to a broker. If the buyer delegate search, the broker holds an auction among a fixed set of sellers to procure the contract and charges a percentage fee for the service. The main primitives of the model are the distributions of buyer search costs and seller service costs.

It is important to discuss an assumption that stands in contrast to previous empirical research on search in the industrial organization literature, which has assumed that sellers' costs are constant across buyers (for instance, Hortaçsu and Syverson (2004) and Hong and Shum (2006)). In models with constant cost, price dispersion is due to variation in search cost under a mixed-strategy pricing equilibrium. While this approach is sensible in a posted-price setting, in which goods are not customer-specific, it is less appropriate in a market for a highly individualized service. The constant-cost assumption has the advantage that search costs are often identified from price and quantity data or even from price data alone. The identification of the model in this paper is complicated by the need to recover two primitive distributions from observed prices in the market − the search cost for buyers and the cost of service provision for sellers. Key to the identification strategy are two features of the setting. First, brokers award contracts through a *competitive bidding process* resembling a first-price auction. Second, I observe both brokered and non-brokered contracts in the data. Carter cost can be identified from broker data following standard arguments in the empirical auction literature. Identifying the distribution of search costs is complicated by the fact that I do not know how many quotes the buyer requests from sellers. The number of competitors is a result of buyers' optimal search strategy, which, in turn, depends on the distribution of search cost. Both the equilibrium number of price inquiries and the search cost are unobserved. Given a distribution of seller costs (which is already identified), the observed variation in prices can be directly mapped to the number of price quotes that buyers must have asked for. The distribution of the number of price inquiries can in turn identify the distribution of search cost for sellers.

My estimates suggest that search costs make up a significant percentage of buyers' total expenses, ranging from about 8% to 15%, and larger volume buyers have larger search costs.

In the main counterfactuals the ability to contract through brokers is removed. This alternative market scenario reveals that both the *direct* and *indirect* effects are large. Expenses for buyers that were using brokers rise, on average, by $445 (11.7%) annually if they have to contract directly through the search market. Prices in the search market rise because the average buyer now compares fewer prices, which reduces the competitive pressure on sellers. Therefore, expenses for buyers that were already searching by themselves rise, on average, by $64 (2.5%).

As a result, the search externality through intermediaries has strong implications for the distribution of rents in the market: while buyers that never use an intermediary benefit less, there are many more of them. Without accounting for the search externality one would underestimate the positive effect of intermediaries on consumer surplus by more than 42%. Taking everything into consideration, I can bound the total annual welfare benefit from intermediation, which lies between \$4.3 and \$12.6 million. The lower bound is a 4.4% increase in welfare and the upper bound a 14.2% increase.[3]

1.1 Literature Overview

This paper relates to the empirical literature on search and intermediation as well as on auctions.

Related Literature on Intermediation: Here, the focus is on intermediaries that help buyers *search* for sellers, but intermediaries also function as guarantors of quality and liquidity and act as market makers (Spulber (1996a)). The theoretical literature has extensively studied the different facets of intermediation (for example, Rubinstein and Wolinsky (1987); Gehrig (1993); Spulber (1996b); Lizzeri (1999); Rust and Hall (2003); Moraga Gonzalez et al. (2014)).

Recently there has been a growth in empirical studies of intermediaries. The most relevant empirical paper is Gavazza (2016), who examines the effect of intermediaries in the secondary market for business aircrafts using a dynamic search and bargaining framework. My paper differs in conceptual ways, as well as in its model, application, and findings. While Gavazza (2016) is interested in intertemporal trading frictions, my focus is on the pricing incentives of a finite number of firms with market power, which are exacerbated by customers' search costs. Thus, Gavazza (2016) tries to explain only aggregate state-dependent prices in the market, while in this setting, the unobserved sources of heterogeneity are crucial for buyers' strategic choice to delegate to an intermediary, as well as for explaining the full distribution of transaction-specific prices. My modeling choices reflect these conceptual differences. Another important distinction is that in my model

[3]Welfare can only be bounded because total change in welfare depends on the fixed cost of broker services, for which I have no estimate.

buyers choose whether to use an intermediary, whereas in Gavazza (2016) buyers are matched exogenously with a seller or an intermediary.

Another paper that compares a bilateral market with an intermediated market is Hendel et al. (2009), which studies the relative performance of a real estate listing service with a platform on which house owners sell their homes directly.

Related Literature on Search Cost: This study is concerned with markets in which consumers lack full information about prices. McCall (1970) and Stigler (1961) were the first to describe buyers optimal search-strategy in a sequential and non-sequential search setting respectively. The corresponding equilibrium models for sequential and non-sequential search settings were formulated in Stahl (1989) and Burdett and Judd (1983). Price dispersion in these models arises due to mixed strategy pricing even though the goods are homogeneous. For the purposes of this study, the assumption of non-sequential search has several advantages. First, under this assumption, the firm's problem against the searching consumer (in the search market) is equivalent to a first-price auction with an unknown number of competitors. This makes the problem tractable and allows me to build on the tools in the empirical auction literature. Second, this assumption also increases the transparency of market comparability and identification. Third, several studies have found that non-sequential search better explains actual search behavior (De los Santos et al. (2012) and Honka and Chintagunta (2014)).

One goal of this study is to quantify the size of search externalities. The theoretical literature has explored such externalities, which arise if firms cannot distinguish between different types of consumers. Two examples are Salop and Stiglitz (1977) and Armstrong (2015). This paper is, to my knowledge, the first empirical to addresses the importance of such an externality and the effects of intermediaries in concentrated search markets more broadly. The goal of most empirical studies in this literature is to back out unobserved customer search costs. Hortaçsu and Syverson (2004) document price dispersion in the mutual fund industry and estimate a search model that allows for product differentiation, using both price and quantity data. Hong and Shum (2006) propose a procedure for estimating search cost from price data alone both for the sequential and non-sequential cases. Allen et al. (2014) use Canadian mortgage data along with quasi-experimental variation due to a merger to estimate search cost non-parametrically. With the exception of Allen et al. (2014), all of the aforementioned papers assume homogeneous cost on the seller side. The fact that these studies explore retail settings, in which consumers purchase from firms, makes this a plausible assumption. In this setting, however, the buyers are firms themselves, and both observed and unobserved variation determines how costly it is for sellers to service the buyer.[4] Thus, I allow cost to be customer-specific. The empirical model, therefore, has two distributions of unobservables: search cost and service cost. The empirical challenge lies in identifying both of these. Another related study is Allen et al. (2019), which estimates a search and bargaining model for the mortgage market, in which home banks have an incumbency advantage.

Related Empirical Literature on Auctions: My model is related to the empirical literature on auctions. Brokers in this market explicitly use procurement auctions to allocate contracts to sellers, but competition in the search market can also be viewed through the lens of competitive bidding: a customer chooses the carter that quotes the lowest price, but carters in that case do not know the number of competing firms, which depends on the search costs of the customer and the optimal search strategy against the known distribution of prices. The pricing sub-game of sellers in the search market can therefore be viewed as a first-price auction with an unknown number of sellers. The identification of auction models has been discussed in Guerre et al. (2000) and further developed in Athey and Haile (2002) for asymmetric auctions.[5]

[4]Factors that determine the cost of service provision include the location, the quantity, the composition of the waste, the distance to the transfer station and many unobserved factors.

[5]The methods have been used to investigate auctions with resale (Haile (2001)), entry into auctions (Li and Zheng (2009)),

Roadmap: Section 2 provides relevant industry facts and describes the data. Section 3 establishes important descriptive facts that inform the setup of the model. Section 4 describes the model, Section 5 the identification of the model, and Section 6 the estimation. Section 7 describes the results, and Section 8 the counterfactual computations.

2 Data and Industry Facts

This section first gives an overview of the data and then establishes two important facts about the market. First, the market supports a large number of suppliers in a geographic area, allowing buyers to choose among many different carters. Second, brokers procure contracts through a *Request for Proposals*, which is akin to a first-price auction. This procurement system will be important in the identification of the model, which I discuss in detail in Section 5.

2.1 Data

Trade waste industry is the official name for New York's private waste market.[6] To free the trade waste industry from its ties to organized crime, Mayor William Louis Giuliani established the *Trade Waste Commission* in 1995. This commission, subsequently renamed the *Business Integrity Commission* (henceforth BIC), has a comprehensive oversight mandate over the trade waste industry.[7]

time incentives in procurement projects (Bajari and Lewis (2009)), collusion in auctions (Asker (2010)), bid preference programs (Krasnokutskaya and Seim (2011)) and many others.

[6]New York's residential waste disposal market is publicly administered by the Department of Sanitation. Self-haulers are an exception. Self-haulers need to register with the BIC to do so. The registration fee for a two-year term is $1000 and $400 for each utilized vehicle; see https://www1.nyc.gov/nycbusiness/description/self-hauler-registration (accessed on 06/14/20).

[7]Private waste carters need to be licensed with the BIC, which monitors their financial and operational activity. The BIC also sets a rate cap for the market and sets rules about sub-contracting and merger applications. If measured in cubic yards this rate cap was $12.2 before 2008, $15.89 from 2008 to 06/2013, which is the relevant data period, and $18.27 from 07/2013 onwards. If measured per 100lbs, it was $8.00 before 2008, $10.41 from 2008 to 06/2013, and $11.98 thereafter.

This paper uses a subset of BIC data, which have been modified to preserve the anonymity of customers and carters. The data cover the period from July 2009 to June 2014 and include all contracts in that period; they contain the customer's zip code, the negotiated price (quoted in terms of either volume or weight) and the quantity of waste generated by the customer. Additional information for each contract includes the date on which it was signed, whether or not the contract was brokered, the type of waste and to which transfer station the waste was carted. In total, there are 1,184,641 panel observations at a half-yearly frequency.[8]

Table (1) Summary statistics

Variable	**Mean**	**Median**	**Std. dev**
Monthly charges ($)	198	92	239.2
Price	11.93	12.2	3.19
Monthly quantity (cu/yd)	25.94	8.66	55.5
Recyclables (yes/no)	0.5	0.0	0.5
Number of weekly pickups	5.27	5.0	3.69

Note: In total there are 1,184,641 observations. The quantity and pickup variables are winsorized at the 1% level to account for outliers.

Table 1 provides some summary statistics. The mean monthly charges for a business are about $198, with a very large standard deviation reflecting the tail of extremely large waste generators. The median number of pick-ups per week is five. Close to half of all businesses generate recyclables.[9] One data caveat is that I only observe the prices that brokers charge customers in 2014, and I cannot match these prices to the customer register contract data. Therefore, I run a hedonic regression of the broker charges on the set of observables that are available in the

Another important restriction regards the length of contracts, which cannot exceed two years, after which the customer has the option to sign with another carter.

[8]This excludes contracts that only involve medical waste, shredding of paper and cardboard or grease haul.

[9]Contracts differ in whether the price is charged per cu/yd or per lbs. For my analysis I only use volume based contracts, which comprise three quarters of my data. Otherwise, I would need to take a stance on the conversion rate between the two, which would introduce a lot of measurement error.

broker dataset as well as in the customer register and impute the broker fee for a given contract as the predicted value from this regression. See Section 6 for details.

2.2 The New York City Market

With very few exceptions, all waste-producing businesses and other private institutions in New York City are required to have a contract with a private waste carters. A breakdown by business type is provided in Table 8 in Appendix D. Carter service consists of the pickup and hauling of recyclable and non-recyclable refuse to one of the 61 transfer stations in or around New York (New York locations are shown as white triangles in Figure 1).[10]

The yearly market volume of the trade waste industry in New York is about $352 million, and it accounts for 3.9 million tons of waste.[11] On average, there are 94 active carters per reporting period (half-year) who serve 110,000 customers. To give a sense of the concentration of the industry: averaged over all reporting periods, the four biggest firms serve 37% of all customers, the seven biggest firms 48% of all customers and the ten biggest firms 55% of all customers.[12]

A salient feature of the New York market is the large number of suppliers serving each geographic area. On average, a zip code is served by 20 carters (Figure 1), which is more than a fifth of the total number of operating firms. The fragmented supply is surprising, especially since carters' services allow relatively little room for horizontal product differentiation.[13] The following observation from an article

[10]The typical volume of one of the rear loaders used in New York City is about 20 to 30 cu/yd, and the median monthly quantity generated by a business in the city is about 8.66 cu/yd. I am not looking at the market for hazardous material, which is subject to much more stringent regulation. Most of the carters operating in New York are not vertically integrated with the transfer station and, therefore, need to individually negotiate tipping fees with the transfer stations.

[11]The publicly administered market creates an additional 3.8 million tons. See Commission (2012). Medical waste and waste shredders are excluded.

[12]For confidentiality reasons, I cannot provide a more complete picture of the firm-size distribution.

[13]There are several potential explanations for the large number of local suppliers, which stands in contrast to the consolidation in other parts of the country. Historically, the waste industry in New York was captured by organized crime (see, for example, Jacobs et al. (2001)). When the waste industry was consolidating on a national scale, New York City was still in

in *The New Yorker* emphasizes this point:

"When I recently walked down a four-block stretch of Broadway on the Upper West Side of Manhattan, I identified about forty businesses — restaurants, clothing shops, bodegas, banks. Licenses in windows listed the commercial-waste haulers they use — at least fourteen in all, by my count, for a stretch that covers only a fifth of a mile. If there was a pattern, I couldn't grasp it: the Starbucks at Ninety-third and Broadway uses a different commercial-waste company from the Starbucks at Ninety-fifth and Broadway." — **The New Yorker, 2009**[14]

The large overlap in suppliers' routes has also received attention from city officials, who have started to collect route information from carters to decide whether to switch to a procurement system with exclusive territories.[15] Prices in the market are individually negotiated between customer and carters. Thus, unlike in most retail settings, searching in this context requires calling an individual carter and haggling. The large number of suppliers and the idiosyncratic nature of the arrangements suggest that this search process is costly for buyers.

the grip of a property-rights and racketeering system. According to Kelly (1999), this system was in place for more than 50 years, and the New York District Attorney's office estimated over-billings of 30% to 40%, which was regarded as a "garbage tax" for doing business. In 1995, a New York City grand jury indicted 23 carters for price fixing, bid rigging, racketeering, corruption and the establishment of a property-rights system. According to the District Attorney's office, small business owners were paying $15,000 a year for the waste removal services and restaurants about $50,000. The article also mentions that while in other parts of the country, many firms were replaced by high-technology entrants, New York was still served by 600 labor-intensive small carters. This might explain why consolidation is delayed in the city and that there is still a large number of relatively small carters. Another potential explanation is the population density of the city. Anecdotally, route density is an important aspect for carters to reduce cost, leading to strong network effects. See also: Nguyen and Wilson (2010). New York City, however, is so densely populated that route density by itself might not be an important margin for overall cost reductions compared to other cities or rural areas.

[14] http://www.newyorker.com/business/currency/a-better-way-to-take-out-the-garbage (accessed on 06/12/2020)

[15] See for example: http://citylimits.org/2015/05/19/city-weighs-reining-in-private-garbage-collectors/ (accessed on 06/12/2020)

Figure (1) Average number of active carters per zip code

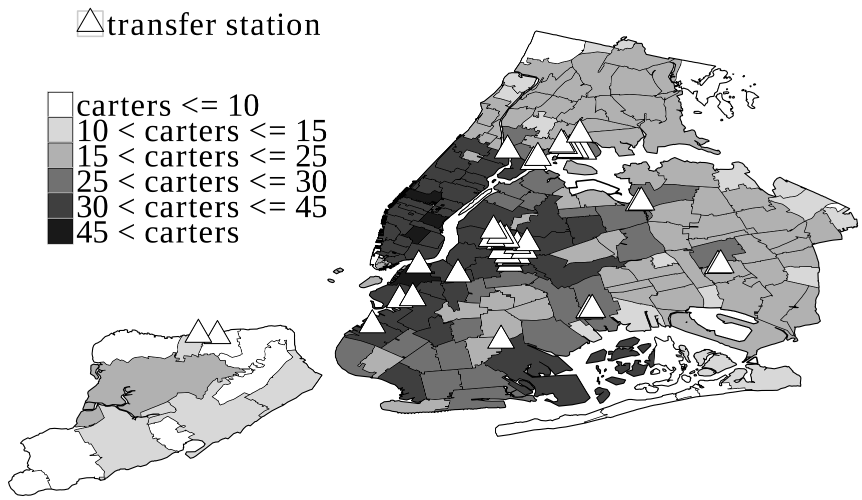

Note: This map shows the number of carters that are active in a zip code, averaged over time. "Active" is defined as having at least one customer in the zip code. Triangles show the locations of the transfer stations.

2.3 Brokers

Trade waste brokers are a potential remedy for customers' search problem, and their services are particularly valuable in markets like New York's. Brokers allow customers, who often operate on a national scale, to have a one-stop shop for dealing with the fragmented landscape of waste removal services. Since many carters in New York City are relatively small, they often are not able to offer the kind of easy access to their services that large waste removal companies can afford. The burden of matchmaking, therefore, falls predominantly on customers. Unlike customers, trade waste brokers know which carters are available to serve customers and have established contacts with a subset of available carters.[16]

Conversations with brokers reveal that they award contracts through a *Request*

[16]Brokers arrange contracts between businesses and carters; they are not allowed to accept direct payments by carters, see http://www.nyc.gov/html/bic/downloads/pdf/regulations/tw_title17_chap_1.pdf (accessed on 08/15/15)

for Proposals - a competitive bidding process akin to a first price-auction - to these carters.[17] A 2015 article in the *New York Times*, which portrays one of the large trade waste brokers, reiterates some of these points.[18]

The identification of the model will build on the fact that the bidding process in a request for proposals works the same as the well understood mechanism of a first-price auction. This bidding mechanism helps pin down the supply-side cost distribution from the subset of brokered contracts. Once the cost of service provision is known, the variation in search cost can be identified from the search market.

The fraction of brokered contracts is relatively stable across business types, locations and volume of buyers. About 13% of businesses indicate that they arranged contracts through a broker. Small waste generators in the first quartile (of quantity of waste) are less likely, at 6%, to use brokers than the three remaining quartiles (15%, 14%, and 19% brokered contracts).[19] Table 9 in Appendix E provides a more detailed view, showing the percentage of brokered contracts conditional on Borough location, the business type, the quantity of waste, and whether a business produces recyclables. The percentage of brokered contracts is similar in Manhattan, Brooklyn and the Bronx (14%, 14% and 17%) but lower in Queens (9%). Large

[17] As part of this research, I spoke to many brokers on the phone. All of them explained that contracts are awarded through request for proposals, which is a competitive bidding procedure.

[18] "*[...] Two big national companies, Waste Management and Republic Services, dominate the market, owning fleets of trucks and hundreds of landfills. Thousands of smaller, regional trash haulers fill in the gaps. Rubicon, based in Atlanta, isn't in the business of hauling waste. It doesn't own a single truck or landfill. [...] It begins by holding an online bidding process for its clients' waste contracts, fostering competition among waste management businesses and bringing down their prices. [...] Through a combination of big data and online auctions for hauling contracts, Rubicon says it reduces clients' waste bills by 20 percent to 30 percent. [...]*" — **The New York Times, 2015**. `http://www.nytimes.com/2014/10/26/business/dividing-and-conquering-the-trash.html?gwh=3238DE6F0B98ED349E78CCC23645724D&gwt=pay&assetType=nyt_now&_r=0` (accessed 07/14/20)

[19] These numbers pertain to the subset of contracts that are used in the study. In Appendix A I explain how I select the sample. Note also that a larger quantity of waste does not necessarily mean a larger business in terms of revenue, but one expects the two to be correlated. There might be multiple reasons for this pattern. It could be due to the differentially stronger discounts for large quantities in the broker market. In the following section, I demonstrate that brokers do price observable information, such as the quantity of waste generated by a business, more systematically. A second reason could be that sales offices need to go to a "procurement procedure" through a broker to comply with audit rules.

hotels and institutions are more likely to use brokers than other non-food retail and wholesale businesses.[20]

Similar to the carter market, the broker market is relatively unconcentrated. The top five firms account for 26.32% of customer market share, the top ten firms for 52.6%, and the largest 15 firms for about 78.95%, averaged over the entire observation period. Note, however, that for the broker market the market shares in New York City might not be the most meaningful. Anecdotally, many brokers are contacted by businesses out of town and might have a more dominant market share in their respective location.

3 Descriptive Results

This section establishes several facts about prices in the trade waste market. First, there is residual dispersion in prices, which points at large expected returns to searching, even when sellers and buyers account for observable information. Second, comparing brokered prices and search-market prices provides evidence that customers with higher search cost use brokers.

3.1 Evidence of Price Dispersion

I document both variation in raw prices p_{ijt} and, following Allen et al. (2014), the dispersion in residual prices to account for observable price variation. Residual prices $\widetilde{p}_{ijt}$ are computed from the following regression $p_{ijt} = \mathbf{X} \cdot \boldsymbol{\beta} + \widetilde{p}_{ijt}$. I restrict the sample to customers that are in the retail business to limit the possibility that observed price variation is driven by unobserved differences across customer

[20] My model assumes that substitution between brokers and the search market is based solely on search cost. But there might be other factors, such as service quality and reliability, that are driving the decision to search through a broker. However, one reason that speaks against this is the fact that virtually all carters that are active in the broker market also serve contracts in the regular market. It is hard to imagine that the workers that serve both types of contracts on a given route know which locations have been brokered and treat them differently.

types. I report the percentage of explained variation and the residual price dispersion from two sets of regressions. The left-hand side variable of these regressions is the price (per cu/yd) that carter j charges customer i during observation period t, which is a half-year. I include only the initial price of each contractual relationship. I run both regressions separately on brokered and non-brokered contracts and include a common set of controls; the only difference is whether I include carter fixed effects. Throughout, I use q to denote the quantity of waste produced by a given customer.

Table (2) Documenting price dispersion

	First specification (no carter FE)		Second Specification (carter FE)	
	Not Brokered	**Brokered**	**Not Brokered**	**Brokered**
$1 - R^2$	0.71	0.35	0.51	0.24
$SD(p_{ijt})$	2.9	4.4	2.9	4.4
$SD(\widetilde{p}_{ijt})$	2.5	2.2	2.13	1.97
$mean(p_{ijt})$	12.4	10.9	12.4	10.9

Note: The table shows the percentage of unexplained variation in observed prices p_{ijt} as well as the dispersion of rates and residual rates $\widetilde{p}_{ijt}$. The second specification is identical to the first but also includes carter fixed effects. All regressions include the following controls: zip code fixed effects, transfer station fixed effects, quantity and higher orders of the quantity variable, recyclables fixed effects, time fixed effects, fixed effects for the number of pickups per week.

Table 2 reports the results of these regressions. Despite the fine-grained controls, the percentage of unexplained price variation is substantial across specifications, and the standard deviation of residual prices is still large. Results from the specification with carter fixed effects show that even within carter, price variation is large. The mean and the percentage of unexplained variation for brokered contracts is smaller, which means that prices on these contracts are better explained by observable differences in contracts.

As a measure of price dispersion, the literature often refers to the coefficient of variation, which divides the sample standard deviation by the sample mean of the price distribution. The coefficient of variation in the search market is 0.24 and in the broker market 0.2. As a comparison, in a well-known empirical study of price dispersion, Sorensen (2000) reports an average coefficient of variation of 0.22 in the retail market for prescription drugs.[21]

3.2 Comparing prices across Broker Market and Search Market

This section establishes that prices including the broker commissions are higher than prices in the bilateral search market, which are, in turn, higher than prices in the broker market without commissions. This relationship calls for an unobserved variable that makes buyers who use broker services differentially inelastic to the prices charged in this part of the market. This pattern is consistent with buyers in the brokered marked having search costs at least as high as the search costs for buyers in the bilateral market.

Figure 2 shows that this clearly holds when examining the raw CDFs of prices. The distribution of brokered prices with commission stochastically dominates the distribution of non-brokered contracts, which, in turn, dominates the distribution of broker prices without commission. Brokers charge substantial markups over the rates obtained from carters in the auction. The median commission is 34.0%, and the mean commission is 38.0%. Results from two sets of regressions (shown in Table 10 and Table 11, in Appendix E) confirm that this relationship also holds after accounting for the observable differences in contracts. The two sets differ only in the dependent variable. The first set includes broker prices without commissions

[21] The following calculations give a sense of the dollar value of this dispersion: for a business that reports in cu/yd and generates the mean quantity of waste (17.8 cu/yd), moving one standard deviation ($SD(\widetilde{p}_{ijt})$) in the residual distribution of prices would cost an extra \$1281.6 over the length of a contract, which is slightly longer than two years. Most businesses in New York City are small. According to the 2013 County Business Patterns provided by the Census Bureau, 60,856 out of 105,439 businesses have fewer than four employees, and 77,965 have fewer than ten employees. BizBuySell.com provides the median sales price of a business in New York City, that have been sold over this platform, which was about \$229,500 in 2013. Assuming an interest rate of 3%, that would imply annual profits of \$6885.0. The subset of businesses sold on this page is almost surely biased towards small companies, but it provides some reference point for the above calculations.

Figure (2) CDF for non-brokered/brokered prices (w and w.o. commission)

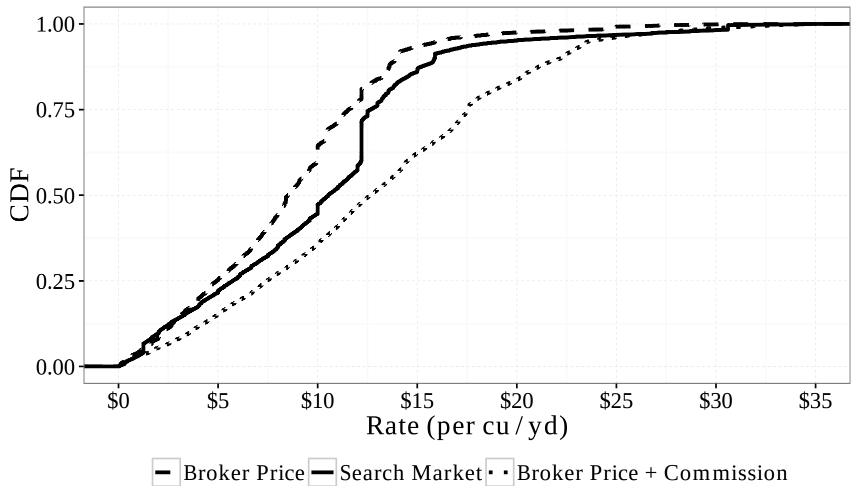

and the second set with commissions. The main variable of interest is a dummy that indicates whether the contract is brokered, and I explore both the mean effect in an OLS specification and the effect at different points of the distribution using quantile regressions. All of the regressions include the following controls: quantity of waste; transfer-station fixed effects; zip code fixed effects; business-type fixed effects; length-of-contract fixed effects; recyclable materials fixed effects; reporting-date fixed effects; number of weekly pickups; and the zip code HHI index.[22] The estimated linear effect of the broker dummy from the first specification in Table 10 is -1.55, which is about 13% of the mean price and shows that brokers obtain lower average prices from carters after observable information is taken into account. Results from the quantile regressions reveal that this difference is composed of a stronger difference in the lower tail of the distribution (25th percentile) with an effect of -1.78. The median effect is -1.0, and the effect is weaker in the upper tail, -0.08 for the 75th percentile. All but the 75th percentile effect are highly significant.

[22] Goldberg (1996), for example, uses quantile regressions to re-examine experimental results on race-based price discrimination. The quantile regressions do not include controls for transfer stations. This is because it is computationally very costly to include that many controls.

Table 11 shows the same set of regressions, where broker prices include the commission. The coefficient on the main independent variable of interest, the indicator whether a contract is brokered, reverses sign from negative to positive. The effect implies that final prices are about \$4.47 higher for brokered contracts, holding other characteristics fixed. At the low end of the distribution, final prices for brokers are \$0.24 higher, while median prices are \$2.6 and prices in the upper tail \$4.7 higher. All of these differences are highly significant. These regression results confirms the observation in the raw price distributions in Figure 2.

Summary of descriptive results In this section, I show that residual variation in prices, even after conditioning on many important observable aspects of a contract is large. This dispersion suggests that customers benefit from searching for carters. Residual dispersion among brokered contracts is lower. Brokers obtain lower prices from carters, but because of high commissions, the average broker price that customers pay is higher. These two facts are consistent with the idea that buyers with high search costs use brokers.

4 Model Overview

To establish a causal effect of intermediaries in the market, one would ideally compare markets with different and randomly assigned prices for broker contracts and test how markets fare under different conditions. However, such exogenous variation does not exist in this setting. To evaluate the effect of brokers, one needs to take into consideration that different observed fractions of brokered contracts might be due to different inherent market characteristics that affect both the propensity to hire a broker and prices. The model will help explain how the observed prices depend on customer search cost and the cost of service provision for carters, as well as account for the potential selection of customers across the two markets and how this selection affects carter pricing.

The search of businesses for brokers is modeled as a sequential game between customers and carters. Brokers are non-strategic players. A customer j in the model is described by privately observed *iid* search cost κ_j, drawn an from a continuous distribution $\mathcal{H}(\cdot)$. These search costs are the marginal cost for an additional price inquiry. The cost κ_j should be thought of as capturing both *search* cost and also the cost of haggling to get the best possible price quote from a given carter. Carters are indexed by i and draw an *iid* customer-specific service cost c_{ij} from a continuous distribution $c_{ij} \sim \mathcal{G}(\cdot)$. In the empirical specification of the model both $\mathcal{H}(\cdot)$ and $\mathcal{G}(\cdot)$ depend on observables. I will suppress the dependence of these key model primitives on observables until I start discussing the estimation since these observables are not essential for understanding the setup of the model and play no role in the identification.

The *timing of the game* is as follows:

> $\mathbf{t = 0}$: Customer draws search cost κ and carters draw service cost c. Both are privately observed.
>
> $\mathbf{t = 1}$: Customers decide whether to delegate search to a broker and, if not, how many price quotes to get, $m \in \{1, \ldots, M\}$, where M is the total number of carters.
>
> $\mathbf{t = 2}$: Carters submit price quotes either in a first-price auction when the contract is procured through a broker, or in the search market.

4.1 Customer Search and Brokers

In the model, I abstract from the competition between brokers. Instead, I assume that buyers make their decision to delegate the search based on the average brokered price (conditional on relevant observables). Brokers determine the seller through a competitive bidding process. I assume that searching customers incur expenses for each additional price inquiry, while I regard the broker infrastructure as fixed. A broker b will, each time she receives a request by a customer, hold an auction with N_b bidders. Let the expected price obtained through a broker be $\mathbb{E}[p^B]$. This price is the average lowest bid over different broker auctions with a varying number of competitors, $\mathbb{E}[p^B] = \sum f_b \cdot \mathbb{E}[p^B|b]$, where f_b is the frequency with which customers encounter broker b, and $\mathbb{E}[p^B|b]$ is the conditional broker price of broker b. In line with the industry practice of running a request for proposal, I model these auctions as a first price-auctions. On the customer side, each type κ will optimally ask for $m(\kappa)$ price quotes, which yields an expected price $\mathbb{E}[p^{1:m(\kappa)}]$. The expression $1 : m(\kappa)$ indicates that the expectation is over the first order statistic out of $m(\kappa)$ draws from the price offer distribution. In the next section, I will provide details on how $\mathbb{E}[p^B]$ and $\mathbb{E}[p^{1:m(\kappa)}]$ are determined in the model. The markup charged by brokers is denoted as ϕ.

With these ingredients I can use Equation 1 to define a cut-off type $\bar{\kappa}$ who is indifferent between an arrangement with a broker and the expected cost of individual search under the optimal search policy,

$$q \cdot \mathbb{E}[p^B] \cdot \phi = q \cdot \mathbb{E}[p^{1:m(\bar{\kappa})}] + m(\bar{\kappa}) \cdot \bar{\kappa}. \tag{1}$$

Every customer with higher search cost will use a broker, and every customer with a lower search cost will search in the bilateral market. Let $\mathcal{F}(p)$ be the equilibrium distribution of price offers in the search market, and remember that $\kappa_j \sim \mathcal{H}(\cdot)$ is the distribution of marginal search cost and $c_{ij} \sim \mathcal{G}(\cdot)$ the distribution of cost draws for carters.[23] I will now discuss what the optimal search strategy of customers against such an equilibrium price offer function looks like. Search is non-sequential.[24] A customer j with search cost κ_j minimizes their total cost over the

[23]The distribution $\mathcal{F}(\cdot)$ is unobserved. The data record only the contract price.

[24]In an earlier stage of this project, I also experimented with a sequential search model. For some parameter values, it was not possible to solve for an equilibrium in these models. This finding, along with the lack of general existence proofs,

number of searches $m \in \{1, \ldots, M\}$:

$$\min_m q_j \cdot \mathbb{E}[p^{1:m}] + m \cdot \kappa_j;$$

using the distribution function of the lowest price in terms of the equilibrium price offer distribution $\mathcal{F}(p)$, the expected cost for making m searches can be expressed as:

$$\min_{m \in \{1,\ldots,M\}} \int_0^{\bar{p}} m \cdot p \cdot q \cdot (1 - \mathcal{F}(p))^{m-1} \cdot f(p) dp + m \cdot \kappa.$$

The following lemma says that buyers, for a given distribution of price offers, will sort themselves according to the optimal number of price inquiries for their type. Depending on $\mathbb{E}[p^B] \cdot \phi$ and prices under the optimal search strategy in the bilateral market, there is a marginal type $\bar{\kappa} < \infty$ such that every type with higher search costs will delegate search to a broker.

Figure (3) Sorting buyers according to search cost

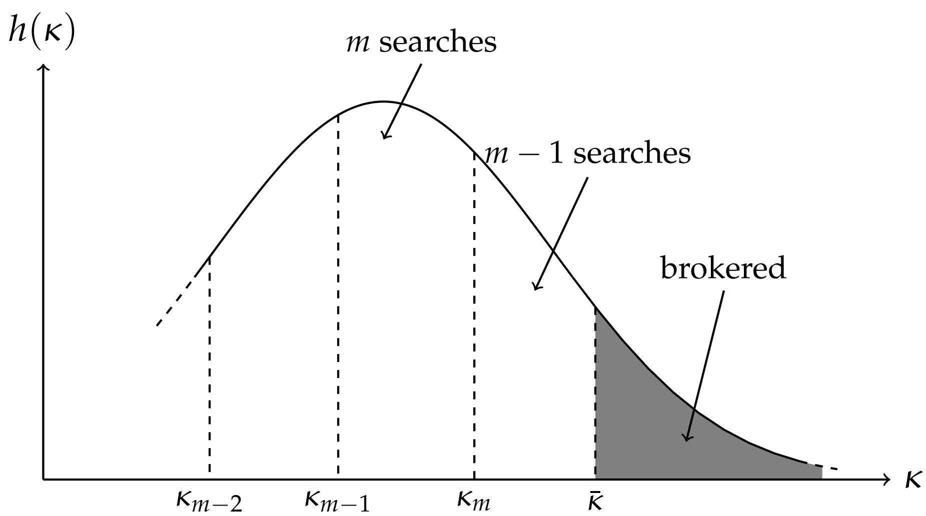

Notes: This figure shows the sorting of buyers according to their search cost. Higher search cost leads to fewer calls m. Types with search cost above $\bar{\kappa}$ delegate their search to a broker.

Lemma 1. *There are marginal types $0 \leq \kappa_{M-1} < \ldots < \kappa_{m-1} < \kappa_m < \bar{\kappa} \leq \infty$ such that every type $\kappa \in [\kappa_{m-1}, \kappa_m]$ samples m firms, and every type larger than $\bar{\kappa}$ delegates search to an intermediary.*

led me to abandon this path.

Proof: This Lemma follows from the fact that $\mathbb{E}[p^{1:m}]$ is concave, and the search cost is linear in the number of searches m. Note that these cut-off types are given by $\kappa_m = q \cdot (\mathbb{E}[p^{1:m}] - \mathbb{E}[p^{1:(m+1)}])$. Because of the sorting it must be true that if $q \cdot \mathbb{E}[p^B] \cdot \phi > \mathbb{E}[p, m] - \kappa \cdot m$ then $q \cdot \mathbb{E}[p^B] \cdot \phi > q \cdot \mathbb{E}[p^{1:(m+1)}] - \kappa \cdot (m+1)$ ■

Brokers provide "rents" to every type above $\bar{\kappa}$. Figure 3 depicts the sorting of buyers along the search cost distribution into bins of types that want to make m searches and the selection into the broker market.

4.2 Carter Pricing

This section describes the symmetric optimal bidding strategy $\beta_b(\cdot)$ of carters in the auction held by broker b, in which they face a known number of competitors N_b. I also describe the symmetric price offer strategy $\beta_S(\cdot)$ in the search market in which carters are bidding against an unknown number of bidders out of M potential bidders (as a function of the buyers search strategy). The search market and and the broker market are linked through the common cost distribution $\mathcal{G}(\cdot)$ that carters in both markets are drawing from. Denote $\widetilde{\mathcal{G}}(\cdot) = 1 - \mathcal{G}(\cdot)$. The optimal bidding functions $\beta_b(\cdot)$ in the procurement auction of broker b with N_b bidders are derived from the following objective function:

$$\max_p (p - c) \cdot \widetilde{\mathcal{G}}(\beta_b^{-1}(p))^{N_b - 1}. \tag{2}$$

In the search market, carters offer their price quotes without knowing consumers' type κ. They serve the contract if they are the cheapest among m firms, where m is a multinomial random variable with probabilities $w_m = \mathcal{H}(\kappa_m | \kappa < \bar{\kappa}) - \mathcal{H}(\kappa_{m-1} | \kappa < \bar{\kappa})$. Note, that the weights w_m themselves depend on carters' equilibrium price offer distribution via the expectations over prices that give rise to the cut-off types. I will now describe carter's strategy $\beta_S(\cdot)$ in the sub-game where carters make price offers to consumers in the search market, planning with the correct vector of search weights $(w_1 \dots w_M)$. These strategies map the customer-specific cost-draw to a price quote. Using this notation, the maximization problem for carters is:

$$\max_p (p-c) \cdot \left[\sum_{m=1}^{M} w_m \cdot \widetilde{\mathcal{G}}(\beta_S^{-1}(p))^{m-1} \right].$$

This maximization problem is akin to a first-price procurement auction with an unknown number of competitors. The number of competing firms is not determined by some entry process but instead by customers' search strategies, as summarized by the weights. Suppressing the dependence on covariates, the bidding function from this problem has a closed form solution (Krishna (2009)) and is given by:

$$\beta_S(c) = \sum_{m=1}^{M} \Big[\frac{w_m \cdot \widetilde{\mathcal{G}}(c)^{(m-1)}}{\sum_{k=1}^{M} w_k \cdot \widetilde{\mathcal{G}}(c)^{(k-1)}} \cdot \Big(c + \frac{1}{\widetilde{\mathcal{G}}(c)^{(m-1)}} \int_c^{\bar{c}} \widetilde{\mathcal{G}}(u)^{(m-1)} du \Big) \Big]. \tag{3}$$

Combining the behavior of carters and the customer search strategy, an equilibrium for the market can be formulated:

Definition 1. *An equilibrium in the decentralized market for customer type is a set of:*

1. *Bidding strategies in each broker market:* $\beta_b(\cdot)$, $b \in \{1, \ldots, B\}$
2. *A bidding strategy in the search market:* $\beta_S(\cdot)$
3. *Customer search weights* $w_1, \ldots, w_M$.

such that:

1. *The search weights result from customer's optimal search behavior under the price distribution* $\mathcal{F}(\cdot)$ *in the search market and* $\mathcal{F}^B(\cdot)$ *in the broker market.*
2. $\beta_b(\cdot)$, $b \in \{1, \ldots, B\}$ *are optimal given the number of bidders* N_b *for broker b and* $\beta_S(\cdot)$, *is optimal given the distribution of price inquiries resulting from* $w_1, \ldots, w_M$.
3. $\mathcal{F}(\cdot)$ *and* $\mathcal{F}^B(\cdot)$ *result from* $\beta_S(\cdot)$ *and* $\beta_b(\cdot)$, $b \in \{1, \ldots, B\}$.

5 Identification

This section discusses the identification of the model's primitives, which are $\mathcal{H}(\cdot)$ and $\mathcal{G}(\cdot)$. For brokered contracts the observables are the broker contract prices p^B, the commission ϕ, the quantity q that consumers contract for and the number of bidders N_b. For search market contracts the observables are the search market price p and the quantity q. In addition, the econometrician observes variables x, which are contract-specific observables. These are not used in the identification and are therefore suppressed until I talk about estimation. Note that in the empirical implementation I do not directly observe the number of bidders but a variable that serves as a very good proxy, which is the number of carters that a broker awards contracts to. This issue is discussed in more detail in section 6.2.2 where I also refer to a robustness check to this assumption.

From Athey and Haile (2002) I adopt the following definition of identification. Define a model as a pair $(\mathbb{F}, \Gamma)$, where $\mathbb{F}$ is a set of joint distributions over the vector of latent random variables, Γ is a collection of mappings $\gamma : \mathbb{F} \to \mathbb{H}$, and $\mathbb{H}$ is the set of all joint distributions over the vector of observable random variables. Implicit in the specification of a model is the assumption that it contains the true $(\mathscr{F}, \gamma)$ generating the observables.

Definition 2. *A model* $(\mathbb{F}, \Gamma)$ *is identified iff for every* $(F, \widehat{F}) \in \mathbb{F}^2$ *and* $(\gamma, \hat{\gamma}) \in \Gamma^2, \gamma(F) = \hat{\gamma}(\widehat{F})$ *implies* $(F, \gamma) = (\widehat{F}, \hat{\gamma})$.[25]

Illustrations in Appendix B.1 demonstrate why, from one price distribution alone, service- and search cost are not separately identified. The distribution of observed prices in the search market is given by Equation 4, and depends both on search costs via the weights and carter costs. The simple intuition for non-identification is that the same distribution of prices can be generated either by a

[25] For a similar definition see Matzkin (2007).

combination of low search cost and high carter cost, or by a combination of high search cost and low carter cost.

$$\mathcal{F}^o(p) = \sum_m w_m \cdot \left(1 - \left[1 - \mathcal{G}(\beta_S^{-1}(p))\right]^m\right). \tag{4}$$

It is therefore crucial that the data record the contracts arranged in both the brokered market and the search market. To make some progress, I make the following assumption:

Assumption 1. *Carter bids in the broker market and in the bilateral market are based on the same cost distribution $\mathcal{G}(\cdot)$ with density $g(.)$, and where $\frac{1-\mathcal{G}(\cdot)}{g(\cdot)}$ is Lipschitz continuous.*[26]

Furthermore, I require the following assumption.

Assumption 2. *A decentralized market equilibrium exists and is unique in the data.*

Under these assumptions one can identify the distribution of carter cost and a partition of the consumer search cost distribution as follows.

Given **Assumption 1**, *Theorem 1* in Guerre et al. (2000) allows for the identification of carter cost based on observing the prices p (winning bids) and N_b in the broker market. A known cost distribution for carters can then be used in conjunction with the distribution of bilaterally negotiated contract prices to identify a partition

[26] In an earlier version of this paper the model accounted for persistent differences of carters by allowing for a finite number of different cost distributions. In Appendix B.3 I discuss the assumptions under the previous setup and briefly outline how one can estimate such a model. Additional details can be found in the old working paper version: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3568571`. I found that accounting for persistent cost differences does not affect the counterfactual results much. One can also use observed prices to directly test for persistent cost differences across carters that serve brokers and those that do not. Table 10 and Table 11 show the results of such a test by including an indicator that is one if at least one contract of the carter is arranged through a broker and zero otherwise. The coefficient on this indicator is not significant in explaining contract price variation in the subset of contracts that I use for my estimation. Lastly, within the symmetric model I have estimated a version in which carters that participate in broker auctions are drawing from a truncated cost distribution $\mathcal{G}(\cdot|\bar{c})$, where $\bar{c}$ is an additional parameter. This is a more ad-hoc way to allow for cost differences across the two markets. I find that the estimation recovers a $\bar{c}$ that is so high that it does not lead to a different cost distribution for carters that participate in broker auctions.

of the search cost distribution. The key intuition is that under Equation 3 one can map from the distribution of winning prices to a unique set of search weights (as illustrated in Figure 3) once carter costs are known and "controlled for". This is the result of the following proposition.

Proposition 1. *If $\mathcal{G}(\cdot)$ is known and **A1** and **A2** hold, the search-weights $(w_1, ..., w_M)$ and cut-off types $(\kappa_1, ..., \kappa_{M-1}, \bar{\kappa})$ are identified from prices p (winning "bids") in the bilateral market.*

Proof: The proof is relegated to Appendix B.2.

The argument involves the following steps: (i) Derive an expression for the inverse bid distribution in terms of quantities that are observed or known, which are $\mathcal{G}(\cdot)$ and $\mathcal{F}^o(\cdot)$ and their densities. This gives rise to an ordinary differential equation that has a unique solution from which I can recover $\beta_S(\cdot)$. (ii) Given that $\mathcal{G}(\cdot)$ and $\beta_S(\cdot)$ are known, the weights are uniquely identified based on an inductive argument. (iii) Given the weights, the full bidding function $\beta_S(\cdot)$ is known, which also identifies the conditional expectations of prices $\mathbb{E}[p^{1:1}], ..., \mathbb{E}[p^{1:M}]$. The cut-off types $\kappa_1, ..., \kappa_{M-1}$ are a direct function of those conditional prices and q, and therefore also identified. Lastly note that all objects in Equation 1 are now known up to $\bar{\kappa}$, which therefore identifies $\bar{\kappa}$. The fraction of brokered contracts, which are all buyer types above $\bar{\kappa}$ are directly observed in the data.

Referring back to Figure 3, with Proposition 1 we can identify the weights corresponding to the partition of the search cost distribution as well as the value of the marginal types $\kappa_{M-1}, ..., \kappa_1, \bar{\kappa}$. Note also that this argument allows for the identification of carter cost and search cost *conditional* on q and whatever other covariates one wants to condition the above argument on. In practice, the partition of the search cost distribution is quite fine. The left panel in Figure 11 (Appendix F) provides an example based on my estimates. Each line in those graphs has a corresponding known empirical CDF that is given by the sum of the search weights up to this point.

Additional Restrictions if q is excluded from $\mathcal{H}(\cdot)$: Because search costs are interpreted as a marginal cost, it is natural to exclude q from the search cost distribution. Carter cost, on the other hand, might naturally depend on q because of scale efficiencies. Since each level of q gives rise to a different set of search weights by moving the cut-off types $\kappa_m = q \cdot (\mathbb{E}[p^{1:m}] - \mathbb{E}[p^{1:(m+1)}])$, the identified partition becomes much finer. The panel on the right of Figure 11 shows this partition for the empirical implementation of the model, which is based on four different quantity levels. However, excluding q from the search cost also leads to a restriction on the substitution to brokers because not only does a higher quantity lead to more search by moving the cut-off types $\kappa_1, ...\kappa_{M-1}$ to the right, but it also moves $\bar{\kappa}$ to the right. If one wants to exclude q from $\mathcal{H}(\cdot)$, the model might be too restrictive in explaining the substitution to brokers across different quantity levels. This issue is discussed in more detail in the next section, where I will also allow the substitution to brokers to depend on q.

6 Estimation

This section explains the details of the estimation. Estimation is conducted via full-solution method of simulated moments. This means that the model is solved each time a new set of parameters is evaluated. I first explain how the moments are constructed given a solution to the model, and then how to solve for equilibria given a parameter guess. Lastly, I give a detailed overview over the full estimation algorithm.

To deal with observable contract heterogeneity I rely on a combination of controlling for observables in the estimation and restricting the sample to a more homogeneous set of contracts. I estimate carter cost conditional on a set of contract covariates, which are denoted as x. The contract covariates x are given by borough fixed effects, the presence of recyclables in the customer's waste stream, and the quantity of waste at the customer location. Including quantity accounts for potential scale efficiencies. For search cost I estimate an unconditional distribution.[27]

Regarding the length of the contract, I rely on the fact that the BIC mandates a maximal contract lock-in of two years. The estimation assumes that when customers search they expect to contract for the full two years. I only model the initial two years and do not try to model the dynamics of contract renewals. Hence, in the estimation I only rely on the initial price in the customer-carter relationship. One issue is that some customers might anticipate that they only remain in contract for a few months and therefore have lower incentives to search. To better fit my assumption of a two-year planning horizon, I only use contracts for customers that stay in the contract for two years. The data shows that 58.4% of all contracts last the full two years, 10% for 1.5 years, 13% for one year, and 17.8% for only six month. I provide an in-depth discussion of dynamic considerations in the customer-carter relationship in section 6.2.1.

I further restrict the estimation to a subset of contracts whose customers operate a retail business to minimize the possibility that price variation is driven by unobserved factors that are idiosyncratic to the customer's industry. The retail category is the largest category in my data and does therefore not restrict the sample too severely. See Table 8 for a breakdown of customer business types. These restrictions leave 44417 contracts on which the model is estimated. Some additional details on the construction of the sample are given in Appendix A.

For the estimation it is convenient to impose parametric restrictions on the distribution of unobservables of the model, which are the search- and carter costs. I assume that both are normal distributions, restricted to the positive domain. The

[27]In other versions of the estimation consumer search cost was also a function of the borough, but this made almost no difference to the estimates.

observables x enter carter cost through a linear index restriction. Under these sample restrictions, the cost-distribution for a contract with observables x is assumed to be:

$$\mathcal{G}(\cdot|x) = \mathcal{N}_{[0,\infty)}(m^c(x), \sigma^c) \text{ where } m(x) = \mu^c + \sum_{k=1}^{4} \gamma_k^{c,Boro} \cdot$$
$$\mathbb{1}\{\text{Contract in Borough k}\} + \gamma^q \cdot q + \gamma^r \cdot \mathbb{1}\{\text{Contract specifies recyclables}\}. \quad (5)$$

The specification for the customer search cost distribution is:[28]

$$\mathcal{H}(\cdot) = \mathcal{N}_{[0,\infty)}(\mu^s, \sigma^s).$$

In the empirical version of the model I also allow the relative attractiveness of the broker option to depend on the customer's quantity of waste. Without this shifter the model would not fit the data well due to the restrictions that changes in the quantity impose on buyers' search strategy and, hence, the type distribution. I illustrate this issue in Appendix D.1. The model predicts that higher quantity buyers search more, which leads all cut-off types, *including the broker-marginal type* to shift rightwards. Holding the search cost-cost distribution fixed, the model would therefore predict that there are too few brokered contracts among high quantity buyers because the incentives to search are always high enough. This is, however, not what I observe in the data. Table 9 shows that a larger fraction of high quantity buyers use brokers.

There are two potential explanations that could rationalize the fact that we do not see "enough" search among high quantity buyers. The first is that search costs depend directly on q and are increasing in q. Alternatively, high quantity buyers are "compensated" in other ways such as additional services from brokers. The first explanation has less appeal because search costs are interpreted as a marginal

[28] I have also tried a version where customer search cost depends on the borough, but this made almost no difference to the results, which is why I decided to drop the extra parameters. The difference in average prices across boroughs is less than $0.50.

cost. The latter is a more plausible explanation, which leads me to assume that the relative attractiveness of brokers also depends linearly on q via ψ as follows:[29]

$$q \cdot \mathbb{E}[p^B|x] \cdot \phi(x) - \psi \cdot q = q \cdot \mathbb{E}[p^{1:m(\bar{\kappa})}|x] + m(\bar{\kappa}(x)) \cdot \bar{\kappa}(x). \tag{6}$$

A final data caveat is that I only have access to broker commissions for one reporting period in 2014, the first time the BIC collected such data. There is no common identifier that would allow me to merge the broker commission data with the main contract-level data. Instead, to bridge the two datasets I run a hedonic pricing regression to impute the broker commissions ϕ based on the two sets of variables that overlap in both datasets, which are the carter charges and the customer zip-code. The model is specified as a linear regression of prices on a fifth-order polynomial in carter charges and zip-code fixed effects. At the contract level, the fit of this model is low with an R^2 of 0.22. However, in this study I do not attempt to explain broker competition and contract-level broker prices. Moreover, price dispersion in broker commissions plays no role in the counterfactual, in which buyers do not have the option to use brokers. Instead, I am interested in selection into the broker market, given the prevailing level of commissions and what those commissions reveals about average search costs in a bin of observables. Buyers in the model make the decision to use brokers based on the average broker price $\mathbb{E}[p^B|x] \cdot \phi(x)$. With an R^2 of 0.87 the imputation model works well in explaining the relevant average commissions $\bar{\phi}(x)$, which is defined conditional on the borough and quartiles of carter charges.

The estimation seeks to recover $\boldsymbol{\theta} = \{\mu^s, \mu^c, \sigma^s, \sigma^c, \boldsymbol{\gamma}^{s,Boro}, \gamma^r, \gamma^q, \psi\}$. I search over those parameters, minimizing the distance between data moments and model-simulated moments. I target the first and second moment of the price distribution as well as the number of brokered contracts, each conditional on x.

The estimation method relies on repeatedly solving the equilibrium for each set

[29] In a previous version of the paper the search cost was conditional on q and under this specification the model suggested that search cost is increasing in q.

of conditioning variables. It would be impossible to solve the equilibrium for each unique value in the support of a continuous covariate. I therefore discretize the quantity q, which is the only continuous variable in x. I divide buyers into four quantity bins based on quartile cut-offs. The product of 4 boroughs, the binary variable that indicates the presence of recyclables, the number of reporting periods, and the four quantity bins results in 160 cells, which means that the model must be solved 160 times for each iteration of the objective function. In an abuse of notation I denote an index for a set from this partition as x. The number of observations in this set is denoted N_x, and the collection of indices corresponding to the collection of contracts in a given set is $\mathcal{A}_x$.

I now explain how I construct the moments for estimation given a set of equilibrium objects $\beta_b(\cdot)\ \forall b$, $\beta_S(\cdot), w_1, ..., w_M, \kappa_1, ..., \kappa_{M-1}, \bar{\kappa}$. For each of the aforementioned cells x I match five different types of moments. These moments are based on the mean and the standard deviation of prices for both brokered and non-brokered contracts as well as the fraction of brokered contracts. Note that the model-derived mean of prices in the search market is given by:

$$\Upsilon_S(\theta, x) = \sum_{m=1}^{M} w_k(x) \cdot \int \beta_S(c|x) \cdot g_{1:m}(c|x;\theta)\ dc,$$

where $g_{1:m}$ refers to the lowest order statistic of carter costs out of m draws. It is much easier to simulate this mean instead of computing it directly. Each simulation draw includes a random draw from the multinomial distribution with weights $\{w_1, ..., w_M\}$ to determine the number of searches and then a corresponding number of cost draws from $\mathcal{G}(\cdot|\theta)$, the lowest of which is denoted as $\underline{c}^k$ and mapped to a price via $p_{s_k} = \beta_S(\underline{c}^k|x;\mathbf{w})$. This leads to a vector of simulated prices $\{p_{s_1}, ..., p_{s_K}\}|x$ of length K. This procedure leads to the following approximation:

$$\Upsilon_S(\theta, x) \approx K^{-1} \cdot \sum_{k=1}^{K} p_{s_k}(\theta, x)$$

Similarly, I can simulate the mean of broker prices by drawing a broker according to f_b and then taking N_b cost draws and mapping the lowest cost draw to a price via $p^b_{s_k} = \beta_b(\underline{c}^k_b; N_b)$:

$$\Upsilon_B(\theta, x) = \sum_b f_b \cdot \int \beta_b(c) \cdot g_{1:N_b}(c|x;\theta)\, dc \approx K^{-1} \cdot \sum_{k=1}^{K} p^b_{s_k}(\theta, x)$$

where f_b is the empirical frequency of contracts with broker b. Simulation is especially advantageous for matching the standard deviations of observed prices, which I can simply compute on the same set of prices from the simulation described above. Lastly, it is easy to compute the number of brokered contracts in closed form as $(1 - \mathcal{H}(\bar{\kappa}(x)|\theta)) \cdot N_x$.

To sum up, I use simulated prices for the broker market and the search market, conditional on x, along with the fraction of brokered contracts implied by the model to build the following estimation moments. The moments for the mean of prices are given by:

$$m_{1,S}(\theta, x) = N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} p_i - K^{-1} \cdot \sum_{k=1}^{K} p_{s_k}(\theta, x), \tag{7}$$

and for the standard deviation of prices by

$$\begin{aligned} m_{2,S}(\theta, x) = {} & \left(N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} \left(p_i - N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} p_i \right)^2 \right)^{0.5} \\ & - \left(K^{-1} \cdot \sum_{k=1}^{K} \left(p_{s_k}(\theta, x) - K^{-1} \cdot \sum_{k=1} p_{s_k}(\theta, x) \right)^2 \right)^{0.5}. \end{aligned} \tag{8}$$

The moments for the brokered market $m_{1,B}(\theta, x)$ and $m_{2,B}(\theta, x)$ are defined analogously using the simulated broker prices $\{p^b_{s_1}, ..., p^b_{s_K}\}|x$. Lastly, the moment for the fraction of brokered contracts is given by

$$m_f(\theta, x) = N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} \mathbb{1}\{\text{brokered}\}_i - (1 - \mathcal{H}(\bar{\kappa}(x))) \cdot N_x. \tag{9}$$

So for each x I have one vector of moments given by

$$m(\theta, x) = \left[m_{1,B}(\theta, x) \;\; m_{1,S}(\theta, x) \;\; m_{2,B}(\theta, x) \;\; m_{2,S}(\theta, x) \;\; m_f(\theta, x) \right]'$$

Denoting $\mathbf{m}(\theta)$ the stacked vector of moments across all bins x, the estimation consists of the following optimization problem:

$$\widehat{\boldsymbol{\theta}}_{MSM} = \text{argmin}_{\theta}\ \mathbf{m}(\boldsymbol{\theta})' \cdot \Omega \cdot \mathbf{m}(\boldsymbol{\theta}) \tag{10}$$

I use the inverse of the observed variance of each moment in the data to weight the moments.[30] I employ a non-parametric bootstrap (Efron and Tibshirani, 1994) to obtain standard errors and 95% confidence intervals for my estimates and the counterfactual numbers. I draw $B = 200$ bootstrap samples, each with the same sample size as the original data. Draws are with replacement from the full set of contracts that are used in the estimation.

6.1 The Estimation Algorithm

For each parameter guess, constructing the moments requires solving for the equilibrium objects $\beta_b(\cdot)\ \forall b$, $\beta_S(\cdot)$, $w_1, ..., w_M, \kappa_1, ..., \kappa_{M-1}, \bar{\kappa}$. In Table 3 I describe the steps that are involved in this process. The inner loop is repeated for each x-bin, but I suppress this dependence to lighten the notation. The equilibrium search-weights are derived by simply updating them iteratively according to steps 1 and 2 in Table 3. Starting from a set of weights with equal probability in each entry, I first obtain a bidding function from which I can compute expected prices. From those expected prices I can compute a set of cut-off types which then give rise to new weights and a new bidding function. Note that at each updating step this procedure also determines $\bar{\kappa}$, which determines the types of consumers that search versus contact brokers. This selection on the consumer side links the two markets. I found that as long as I rule out the Diamond paradox — an equilibrium in which nobody searches and all sellers charge the monopoly price — this iterative updating quickly converges to a unique vector after a few iterations.[31] This drastically

[30] For each moment is use 20 simulation draws, which leads to a total of 1280 simulation draws.

[31] The Diamond paradox is rejected by the data.

speeds up the inner loop because I do not have to run a separate numerical search for the equilibrium.

Table (3) Estimation Algorithm

Algorithm 1: Estimation of the Model

Result: Estimate of θ

while $\mathbf{m}(\theta)' \cdot \hat{\Omega} \cdot \mathbf{m}(\theta) >$ `outer tolerance` **do**

1. Initialize weight vector w_0 with $1/M$ in each entry;

2. while $d(w^k, w^{k-1}) = ||w^k - w^{k-1}|| >$ `inner tolerance` **do**

2.1 Recompute the broker bidding functions $\beta_b^k(\cdot|\boldsymbol{\theta})\forall b$;

2.2 Use w^k to recompute the bidding function $\beta_S^k(\cdot|\boldsymbol{w}^k;\boldsymbol{\theta})$;

2.3 Use $\beta_S^k(\cdot|\boldsymbol{w}^k;\boldsymbol{\theta})$ to compute expected prices $E[p^{1:m}]$ for each m;

2.4 Recompute $\kappa_m = E[p^{1:m}] - E[p^{1:(m+1)}]\ \ \forall m \in 1, ..., M$ as well as $\bar{\kappa}$;

2.5 Form new weights $w_m^{k+1} = \mathcal{H}(\kappa_m|\kappa < \bar{\kappa}) - \mathcal{H}(\kappa_{m-1}|\kappa < \bar{\kappa})\ \forall m$;

end

3. Use the equilibrium objects to simulate $\{p_{s_1}, ..., p_{s_k}\}$ and $\{p^b_{s_1}, ..., p^b_{s_k}\}$, compute the average and the standard deviation of simulated prices, compute fraction of brokered contracts $(1 - \mathcal{H}(\bar{\kappa})) \cdot N.$;

4. Construct moments for the objective function.

end

One potential concern is that the set of equilibria with a non-degenerate price distribution is not unique. I have not found this to be an issue in practice. One advantage in this setting is that the equilibrium in the search market is easily summarized — it is a set of search weights. A given set of search weights will automatically imply a bidding function for carters, which is unique given the set of weights. One can therefore easily explore uniqueness by starting the iterative procedure to solve for equilibria from different initial weights. For the parametric implementation that I use in this paper I have never encountered an issue of multiple equilibria. In Appendix G I shows the results from these explorations. These tables show each step of the iteration until convergence; for a wide variety of starting values, the algorithm converges to the same equilibrium.

6.2 Limitations and Practical Considerations for Estimation

6.2.1 Contract Dynamics

The estimation focuses on new contracts between buyers and carters. The welfare quantifications that I discuss below only pertain to the initial contract period. This allows me to avoid some of the complications that would arise from modeling a dynamic relationship once the initial contract period ends. One remaining complication is that, while the data provides information about the duration of a customer-carter relationship, it does not speak to when contract terms are renewed. Thus, one must make an assumption about the length of the initial contract period. One thing that helps in this regards is that per BIC regulation customers can switch after two years at no additional pecuniary cost.[32] I therefore assume that each first time relationship lasts for two years. In addition, to rule out unobserved heterogeneity, I only include contracts in the estimation that lasted for two full years and where the customer did not either exit or switch carters.[33]

A closely related issue is switching cost, which could affect buyers' decisions once they initialize a relationship with a carter. For example, if buyers can go

[32]http://www.nyc.gov/html/bic/html/trade_waste/customer_info_contracts.shtml (last accessed on 08/19/15). In line with this regulation being binding for most contracts I find that about 2.18 years pass, on average, before customers obtain a new rate, go out of business or switch.

[33]Overall, this issue is less problematic than it sounds. Similar to q, contract-length scales the cut-off types and, as a result, the estimated search cost. Suppose that ρ is the number of time periods for which the contract is signed. The cut-off types κ and, therefore, the estimated search cost and contract expenses are both scaled by ρ as the length of the contract changes. This means that search cost per unit of time and as a fraction of total expenses (contract expenses plus search cost) remains the same because the ratio of search cost and total expenses remains the same, as both are scaled by the same factor.

back to the incumbent seller's price but have to search for new outside offers, search costs act as switching costs. If buyers obtain a low initial price quote, this might help them sustain lower future prices in the relationship. For estimation, this would mean that search costs are underestimated.[34] To give a sense of how large the dynamic component of pricing is, Figure 9 in Appendix C shows the time dummies of a hedonic price regression (using the same controls as similar regressions referred to earlier) with and without contract fixed effects. We see that after two years, prices adjust upwards, consistent with search costs acting as switching cost. However, the magnitudes, at less than \$.50, are modest. In general, search and switching costs are hard to identify separately without observing the information set of buyers (Hortaçsu and Syverson (2004)). Embedding the current framework in a dynamic setting that accounts for such dynamics is beyond the scope of this paper and left for future research.

6.2.2 Number of Bidders in The Broker Market

I do not directly see the number of bidders in the broker auctions on an auction-by-auction basis. I therefore set the number of bidders on a broker contract to the number of carters that have won a contract through this broker. Note that this could underestimate the true number of bidders if, in reality, brokers do not solicit quotes from every carter they have a relationship with. Welfare changes in the counterfactual depend on both search and service costs, and a bias in the number of bidders would lead to a bias in the relative size of search and service costs and, possibly, in the conclusions drawn from counterfactuals. If the true number of bidders is smaller than assumed, service costs would be overestimated since the observed first-order statistic is less selected. This, in turn, would mean that search costs are underestimated since sellers are drawing from lower cost distributions;

[34] In that case, one would have $\kappa_m = q \cdot (\mathbb{E}[p^{(m)}] - \mathbb{E}[p^{(m+1)}]) + (V(p^{(m)}) - V(p^{(m+1)}))$, where the second term would be omitted.

thus, it must be that buyers searched less to rationalize prices in the search market. Since the overall effect on welfare is unclear, the counterfactual results section discusses a robustness check, which shows how results change if the number of bidders was, in fact, lower or higher.

6.2.3 Contract Level Cost Heterogeneity

The model does not take into account the possibility of contract-level heterogeneity that is unobserved to the econometrician. The auction literature has established methods for dealing with auction-level heterogeneity that is *observed to bidders* but unobserved to the econometrician. These methods build on insights from the nonlinear measurement error literature (Li and Vuong, 1998) and have been applied in a variety of settings.[35] The key requirement for these approaches is multiple "measurements", which in the auction context would mean multiple bids per auction. In this setting, I only observe what can be interpreted as the equivalent of the winning bid, which makes such methods infeasible. Alternatively, auction level heterogeneity could be unobserved to the econometrician and only *imperfectly observed to bidders* who in this case receive a noisy signal about a *common value*. Such models are typically not identified (Athey and Haile, 2002), but one could test for common values in first price sealed-bid auctions (Haile et al., 2003).

While it is outside the scope of this paper to extent the model in these directions, I would like to briefly discuss the likely bias that would result from such auction-level heterogeneity. In both cases the likely bias is similar and would result from an overestimate of idiosyncratic cost differences across sellers. One welfare effect of intermediaries is higher allocative efficiency. Such allocative effects arise because buyers do not necessarily find the lowest cost seller. The more dispersed the idiosyncratic cost of carters, the more severe is mis-allocation. This means that any allocative effects due to cost differences in the counterfactual are overstated if some

[35] See, for example, Asker (2010) and Krasnokutskaya (2011)

of the variation is in reality due to auction-level variation. Moreover, lower cost dispersion means that the bidder's incremental gain from searching is in reality (unobservably) lower. This would mean that search cost are underestimated. Because brokers increase allocative cost efficiencies less when cost are less dispersed but reduce search cost more when search cost are higher, the overall bias on the counterfactual results is unclear.

7 Results

This section describes the estimation results. I graphs the model's fit in Appendix J. The parsimonious model fits the key moments quite well across the different quantity bins, with broker prices being slightly underestimated for low quantity buyers. Table 4 shows the parameter estimates along with standard errors. Average carter cost per cu/yd are estimated as $9.969 with a standard deviation of 2.96. Note that this is the value for the population distribution and not the average cost for observed contracts, which is a selected set from this distribution. Carter costs are decreasing in the quantity of waste but not very strongly. This means that the strong quantity discounts that I observe in the data are primarily driven by customers' increased incentives to search. The model recovers a positive effect of recyclables on cost but the coefficient is not significant. There are small cost differences across boroughs. The Bronx and Manhattan are estimated to be cheaper to service than Brooklyn and Queens.

Table (4) Model Parameter Estimates

Parameter Name	Parameter Estimate	SE	95% **Confidence Interval**
Supply			
Mean Carter Cost	9.969	0.203	[9.774,10.024]
SD Carter Cost	2.96	0.151	[2.806,3.106]
Quantity cost efficiencies ($\times$1000)	-2.456	1.545	[-4.602,-0.164]
Recyclables cost efficiencies ($\times$1000)	-0.152	6.879	[-1.783,19.157]
Bronx cost shifter	-0.235	0.034	[-0.3,-0.211]
Brooklyn cost shifter	-0.002	0.008	[-0.002,0.021]
Manhattan cost shifter	-0.32	0.035	[-0.407,-0.279]
Demand			
Mean Search Cost	79.718	5.416	[77.178,96.963]
SD Search Cost	62.352	4.613	[58.717,72.486]
Quantity Broker Shift	-0.309	0.015	[-0.334,-0.286]

Note: This table shows the parameter estimates along with bootstrapped (standard errors) and [95% confidence intervals], based on 200 bootstrap iterations. The borough cost shifters are relative to Queens.

The mean search cost estimate is \$79.718 with a large standard deviation of 62.352. The coefficient that captures differential selection to brokers due to differences in quantity is estimated as -0.309. This means that with each additional cu/yd customers behave as if the total broker bill is reduced by about 31 cents. While this effect does not matter for small quantity buyers it can be large for buyers with high quantity and explains why I do not necessarily observe them searching by themselves despite their steep incentives to do so.

Table 5 shows the average cost – per cu/yd – for the contracts that are served (i.e., the cost of "winning" carters) conditional on the borough and buyers' quantity. Contracts with higher quantity for buyers have lower cost. This is mostly because buyers with higher q have higher incentives to search, leading the cost to be more selected. Figure 4 shows this effect graphically. $Q_{a,b}$ in Table 5 refers to quantities from quantile a to b.

Table (5) Expected $-cost for carter per cu/yd per month

	Bronx	**Brooklyn**	**Manhattan**	**Queens**
$Q_{0,25}$	9.8	11.9	9.5	9.3
	(0.68)	(0.665)	(0.697)	(0.694)
$Q_{25,50}$	8.9	9.2	9.5	9.2
	(0.557)	(0.746)	(0.626)	(0.627)
$Q_{50,75}$	7.9	6.9	7.8	7.9
	(0.572)	(0.666)	(0.611)	(0.709)
$Q_{75,100}$	5.4	5.7	5.5	5.7
	(0.334)	(0.415)	(0.422)	(0.376)

Note: The table shows the average cost for different boroughs. To compute this cost I draw repeatedly (10000 times) from the distribution of bidders. $Q_{a,b}$ refers to quantities from quantile a to b. Bootstrapped standard errors are provided in parentheses (200 iterations).

Figure (4) Carter cost

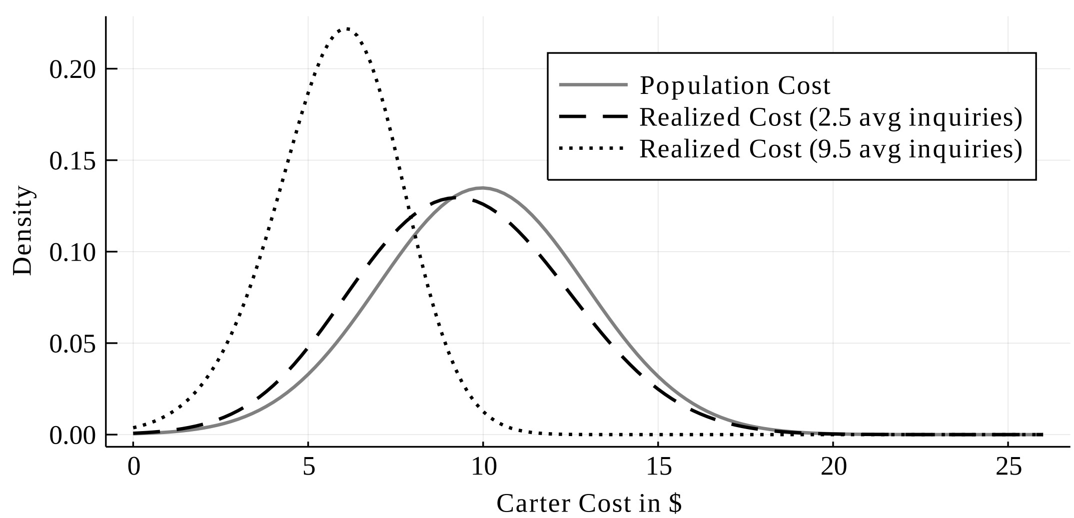

Notes: This figure shows the population distribution of carter cost and the realized distribution for two different equilibrium search weights, one with 2.5 and one with 9.5 inquiries on average.

Table 6 gives a broad overview of customers' search costs. Buyers who contract for higher quantity search more and therefore spend more on searching. However, Table (6) Search Cost and Inquiries

Subset	Search Cost Per Inquiry ($) $\kappa < \bar{\kappa}$ (searches)	$\kappa > \bar{\kappa}$ (brokers)	Number of Searches $\kappa < \bar{\kappa}$ (searches)	Total Search Cost ($) $\kappa < \bar{\kappa}$ (searches)	Fraction of Total Expenses
Bronx					
$Q_{0,25}$	87.3 (23.865)	220.1 (60.25)	1.2 (0.33)	93.0 (25.399)	0.15
$Q_{25,50}$	81.5 (22.347)	194.7 (53.467)	1.6 (0.434)	99.1 (27.071)	0.08
$Q_{50,75}$	83.2 (22.875)	200.7 (55.511)	2.9 (0.794)	163.0 (44.487)	0.08
$Q_{75,100}$	72.1 (19.815)	170.3 (46.798)	8.2 (2.248)	536.5 (146.565)	0.09
Brooklyn					
$Q_{0,25}$	85.6 (23.36)	210.7 (57.478)	1.2 (0.331)	92.9 (25.372)	0.15
$Q_{25,50}$	75.2 (20.491)	177.1 (48.274)	1.7 (0.452)	98.4 (26.86)	0.08
$Q_{50,75}$	57.4 (16.229)	145.3 (40.462)	3.8 (1.04)	165.7 (45.244)	0.08
$Q_{75,100}$	63.3 (17.237)	154.2 (41.995)	8.7 (2.378)	510.9 (139.083)	0.08
Manhattan					
$Q_{0,25}$	87.8 (24.01)	223.6 (61.208)	1.2 (0.33)	93.0 (25.406)	0.15
$Q_{25,50}$	83.4 (22.856)	201.5 (55.336)	1.6 (0.43)	99.3 (27.119)	0.08
$Q_{50,75}$	85.8 (23.502)	212.0 (58.309)	2.8 (0.778)	161.9 (44.192)	0.08
$Q_{75,100}$	75.8 (20.764)	178.6 (48.977)	8.0 (2.199)	545.8 (148.95)	0.09
Queens					
$Q_{0,25}$	85.6 (23.368)	210.6 (57.514)	1.2 (0.331)	92.9 (25.372)	0.15
$Q_{25,50}$	75.1 (20.517)	177.0 (48.335)	1.7 (0.452)	98.4 (26.864)	0.08
$Q_{50,75}$	57.4 (16.027)	145.3 (40.175)	3.8 (1.045)	165.7 (45.251)	0.08
$Q_{75,100}$	63.2 (17.29)	154.0 (42.081)	8.7 (2.375)	510.6 (139.278)	0.08

Note: This table shows expected search cost per inquiry, the number if inquiries as well as total expenses for search. Bootstrapped standard errors are provided in parentheses.

as a fraction of the total expenses, which include contract expenses, their expenses on search are lower.

A separate look at the average search cost for customers that use brokers and those that do not gives a sense of the selection into broker services. The search costs for customers that use brokers are about two to three times as high as those who contact carters directly. The estimates, therefore, reflect the high mark-ups charged by brokers. The wedge that this markup generates between the broker market and the non-brokered market is rationalized through higher search costs for buyers that use broker services. For buyers with higher quantity, the equilibrium cut-off type moves to the left. That is why the average cost per inquiry is lower both above and below the cutoff for higher quantity buyers.

Table 6 also gives an overview of the total costs that customers incur for searching – i.e., their cost per inquiry times the equilibrium number of price solicitations. They range from $93 for low-quantity waste generators up to $545 for high-quantity waste generators, with an average number of inferred price inquiries of about 3.5. To put these numbers into perspective: search costs are between 8% and 15% of buyers' total expenses (contract expenses + search expenses).

The magnitude of the search cost is important for the city's current consideration to move to a procurement system with exclusive territories. Such a mechanism would force buyers to contract with one carter and, therefore, removes the ability to search, but also the cost of searching. A full welfare evaluation should take into account the search cost.[36]

[36] I have also made an effort to compare my estimates to the literature. In Allen et al (2019) the authors compare the search cost from different models (Honka (2014), Hortaçsu and Syverson (2004), Hong and Shum (2006)). To make the estimates comparable across the different settings, they use the ratio of search cost to the standard deviation of payments. I build on this comparison and compute the same ratio. Hortaçsu and Syverson (2004), at a median search cost of 5 basis points, the ratio is 8%. Hong and Shum (2006) estimate an average search cost of $1.58 for the non-sequential search model, yielding a ratio of 33%. Honka (2014) estimates the cost of searching for policies to be $28 per online search and $100 per offline search in the auto insurance market. Depending on which of those estimates is used, her ratio range between 10% for online transactions, and 35% for offline. In my setting I find a median ration of 26% and a mean of 34%. This is broadly in line with the aforementioned papers.

8 Counterfactual Market without Intermediaries

This section discusses the key counterfactual, which explores what happens to the market if buyers cannot make use of intermediaries. The absence of broker services introduces several changes on both sides of the market. On the buyer side, it changes the composition of search costs in the search market, as well as the market price. This raises the overall expenses for the contract. It also increases the number of buyers that have to search. Brokers lose commissions, and carters' expected profits change because buyers in the search market search less. Welfare declines because of larger expenses for search and a higher average contract cost.

Figure 5 and Figure 6 provide an example (based on the estimates for *Manhattan* and the highest quantity quartile) of the changes in the counterfactual and illustrate several important points. First, Figure 5 shows a typical partition of the search cost distribution into buyers who take different numbers of draws from the price distribution. The solid lines are all the marginal κ-types who are buying in the search market when intermediation is still an option. The graph shows that there are two separate effects that decrease the amount of search. The first is a selection effect — the search market in the counterfactual is populated by many more buyers with high search cost. The new buyers entering the search market are all the types above the cut-off type ("New marginal type"). The second is a response to a changing price distribution and slightly more subtle. High search cost buyers entering the market leads carters to raise prices. Importantly, low cost carters raise their margins more than high cost carters, which flattens the bid function. A flatter bid function means that the returns to searching go down, which leads those buyers that were already in the search market to also search less. This effect compounds the selection effect. In Figure 6, this response shows up as marginal types moving to the left. The black dashed line ("Broker marginal") shows the type $\bar{\kappa}$, who is indifferent between using a broker and searching on her own.

Figure (5) Cut-off types and counterfactual cut-off types

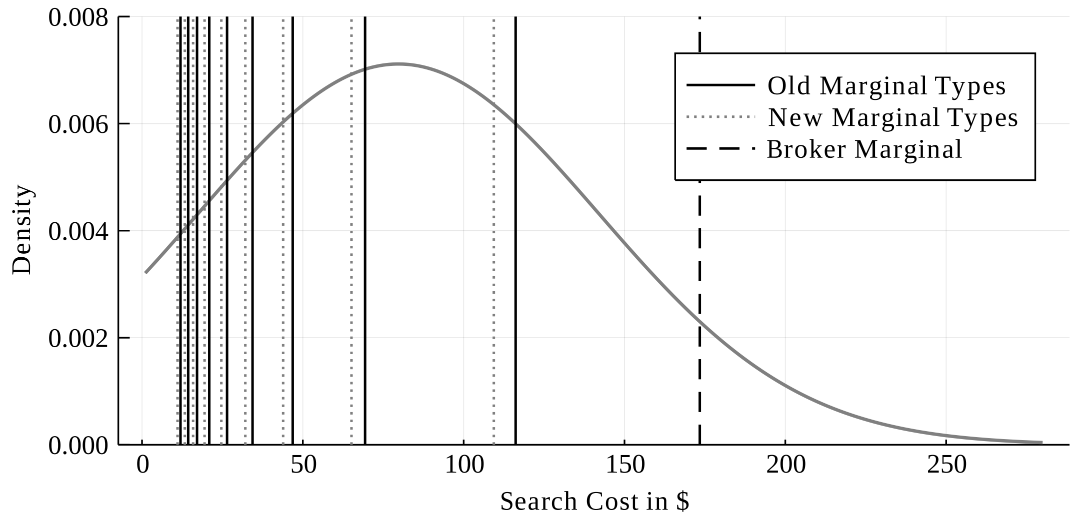

Notes: This figure shows an example of a search cost density along with the old and new marginal κ types.

Figure 6 compares the histogram of price inquiries, with and without brokers, and shows that everybody to the left of the "Broker marginal" falls either into the bin of one price inquiry. The old CDF of price inquiries strictly dominates the new one meaning that there is strictly less search, which leads sellers to increase their price.

Table 7 provides the main overview of the changes in buyer and seller welfare and in overall welfare. Figure 12 in Appendix E breaks the changes down for the different estimation bins. Carters charge, on average, 4% higher prices and their margin slightly increases. However, their overall profits decrease. Carter profits decrease because high quantity buyers (in the highest quartile of quantity) have incentives to use brokers that are unrelated to search cost, as measured by ψ. These buyers have strong incentives to search in the counterfactual, leading to more price draws than brokers are soliciting through the auctions. Appendix E shows the carter profit decrease is coming mainly from buyers with very high quantity, who are responsible for an disproportionate percentage of profits.

Figure (6) CDF of search of equilibrium search strategies

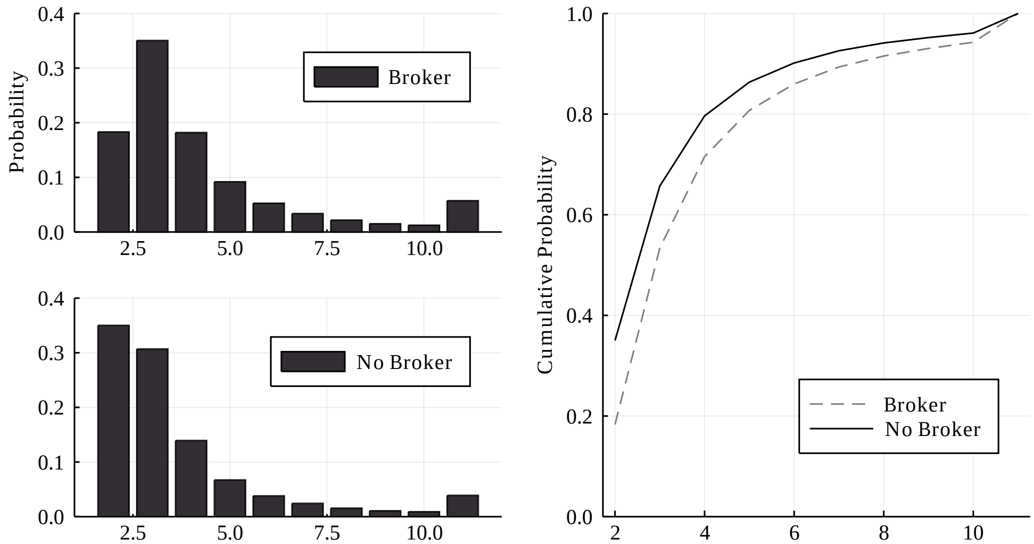

Notes: The right panel shows the cumulative CDFs for the two original and the counterfactual search strategy and the left panel the two histograms.

Buyer expenses include both the contract and search expenses. Eliminating brokers for buyers that were using their services means an increase in expenses of \$445, or about 11.7%. For buyers who were not using brokers prices rise by 4.46%, and their overall cost increases by \$64 or 2.5%. While this indirect effect is smaller, it is important to note that the number of buyers that benefit from the externality in the current market setting is much larger than the number of buyers that use intermediaries. If we multiply by the number of respective buyers, the implied loss for buyers who were using brokers is \$3.2 million, while for those who were contracting directly with sellers the total loss is nearly as large as \$2.4 million. Brokers, therefore, redistribute rents to a large extent through the externality and I would miss 42% of their positive effect on consumer surplus if I did not account for the externality.

Regarding the overall welfare comparison, I need to deal with the fact that the model does not produce an estimate for brokers' fixed cost. But it is possible to

Table (7) Counterfactual Overview

	Change in Buyer Expenses			Carter Margin		Welfare	Welfare
	Not Brokered	Brokered	All	(Not Quantity Weighted)	Carter Profits	Total Cost Lower Bound	Total Cost Upper Bound
Δ Absolute	$64.0	$445.0	$127.0	$0.046	$-11.1	$4.28	$12.61
SE	(11.0)	(150.0)	(30.0)	(0.0285)	(10.0)	(1.06)	(2.91)
95% CI	[48.2,80.4]	[301.7,770.4]	[92.9,189.4]	[0.004,0.087]	[-29.5,0.3]	[3.37,6.81]	[9.98,19.22]
Δ Percent	2.52%	11.7%	4.6%	1.95%	-1.8%	4.41%	14.22%
SE	(0.55)	(1.91)	(1.12)	(1.19)	(1.59)	(1.16)	(3.88)
95% CI	[1.85,3.57]	[9.1,15.38]	[3.35,6.99]	[0.15,3.81]	[-4.63,0.04]	[3.44,7.25]	[11.01,23.3]

Note: This table shows expected search cost per inquiry, the number of inquiries, as well as total expenses for search. Search cost changes are computed under the assumption that brokers' total variable profits are equal to their fixed cost, which provides a lower bound on the change. Bootstrapped (standard errors) and [confidence intervals] based on 200 iterations.

bound these fixed costs and, therefore, the welfare change. The upper bound (on the cost increase - i.e., welfare decrease) can be obtained by assuming that the fixed costs are zero and the lower bound under the assumption that the total fixed costs are equal to the total observed variable profits, which equals the total commission payments. The formulas for the computation of the welfare changes are provided in Appendix I. According to these calculations, the upper bound on the welfare change is $12.6 Million and the lower bound $4.3 Million. At the lower end this would imply a 4.4% decrease in welfare, and at the higher bound a 14.2% decrease.

The welfare loss is a combination of higher search cost and a re-allocation effect. The re-location effect means that without brokers the market engages in less search and picks carters that are less well suited for customers, which results in a higher average cost. The size of both effects varies depending on the quantity of waste for buyers but on average the realized cost per cu/yd is 2.3% higher without brokers.

8.1 Robustness to Number of Bidders in the Auction

One data limitation is that is that I do not directly observe the bids in broker's requests for proposals. Instead, I use the number of broker contacts to infer the number of bidders. This is likely overestimating the number of bidders if not all those carters participate in each auction. However, it could also underestimate the number of bidders if there are additional bidders that bid but never win. To probe the sensitivity of the counterfactual results I re-estimate the model and run counterfactuals under the assumption that the number of bidders is, instead, one third lower or one third higher. The results from this robustness check are shown in Table 18.

Under the more likely scenario where the number of bidders is underestimated, the qualitative and quantitative conclusions are almost unchanged. The externality of brokers leads to a slightly higher welfare loss for non-brokered customers under the counterfactual, rising from 2.5% to 2.9%. For formerly brokered customers expenses rise by 12.1% instead of 11.7%. The total welfare change is now bounded between 4.6% and 14.4% instead of 4.4% and 14.2%. The largest difference is in the outcome of carters, whose margin increases by 0.1% instead of 2.0%.

Under the less likely scenario that the bids are overestimated by 30%, the changes in outcomes would be larger. While the welfare benefits to buyers would be almost unchanged, lowered from 4.6% to 4.5%, the total welfare gain would be larger and bounded between 9.3% and almost 20%. Carters' margin would go up by 5.1% instead of only 2%.

What the results from this robustness check show is that the main conclusions from the counterfactual would be qualitatively unchanged and while there are changes in quantitative magnitudes, those changes are modest.

8.2 Summary and Discussion of Counterfactual Results

Intermediaries in this market redistribute rents from sellers to buyers. By keeping high-search cost buyers away from the search market, they make it more competitive for sellers. Buyers in the search market benefit from lower prices and lower search expenses. Since there are more buyers in the search market, the transfer of rents to buyers is predominantly due to the externality. In general, these services are therefore under-provided since intermediaries create social returns that are not reflected in their private pay-offs. While this result provides a rationale for policies that promote intermediation, it also highlights an important insight for regulatory information provision and disclosure policies (for example, see Jin and Leslie (2003)): it means that such policies can have large effects due to supply-side responses, even if they only reach a fraction of customers.

A natural question is whether these results apply to other markets with search frictions and buyer-specific costs, a common feature of business-to-business markets. Quantitative results will necessarily depend on the exact primitives that govern behavior in other markets. However, one would expect that in many markets that fall within this paradigm, search frictions are even more severe and, thus, increase the welfare effect created by intermediaries. In markets for investment goods, such as construction equipment or production lines, products are often ordered with extremely idiosyncratic specifications, and sellers in these markets are often scattered over many countries, which increases the costs of search and negotiations.[37]

9 Conclusion

This paper studies the competitive and welfare effects of intermediation in a decentralized market. Such a structure is common in retail services markets, wholesale trade markets, and markets for investment goods. Intermediaries can give rise to search externalities that also benefit buyers who do *not* contract through intermediaries. The self-selection of buyers with high search costs into the intermediated market changes the composition of buyers in the search market, thereby making

[37]Note, however, that in markets with extremely idiosyncratic features buyers typically organize their own "request for proposals". This is, for example, true for large architectural projects and capital equipment purchases.

it more competitive for sellers. To quantify these effects of intermediation, I use a new and detailed dataset from the New York City trade waste industry, which provides a comprehensive and rare insight into a decentralized market.

Methodologically, this study contributes to the literature by introducing a new search model that takes into account the idiosyncratic cost of servicing buyers. This model combines elements from the empirical search and procurement-auction literatures. The identification of the model is challenging, since one needs to distinguish between the distribution of sellers' service-costs and buyers' search costs, both of which are unobserved. This stands in contrast to previous empirical studies in the search literature, which estimate only a distribution of search costs under a uniform cost assumption and posted prices. The joint data from the brokered and the bilateral search markets together identify search and service costs. An important institutional detail is that prices in the brokered market are formed according to competitive bidding. This pins down the costs for carters. Known cost distributions can then be used in conjunction with prices in the search market to back out the distribution of search inquires. The latter is then used to recover buyers' search cost.

The estimates reveal that buyers' search costs are an important factor to take into account when designing policy, such as merger guidelines or the trade-off between the existing decentralized market and a system of procurement with exclusive territories, which New York City is contemplating.

Counterfactuals show that intermediaries redistribute a sizable portion of rents from sellers to buyers and improve overall welfare by reducing search costs and by reallocating contracts to lower-cost suppliers. The results highlight the importance of positive search externalities created by intermediaries.

Several possible follow-up projects emerge from the analysis. The model currently abstracts from the competition between brokers, which might be important to consider for some applications, such as finding the optimal subsidy to brokers.

To find the optimal subsidy we would need to know how brokers pass on subsidy payments and, hence, how they compete. Another possible direction for further research is to extend the model to allow for differentiation in the quality of service provision or the dynamic aspects of buyers' choice among sellers.

References

Allen, Jason, Robert Clark, and Jean-François Houde, "The effect of mergers in search markets: Evidence from the Canadian mortgage industry," *American Economic Review*, 2014, *104* (10), 3365–96.

_ , _ , **and** _ , "Search frictions and market power in negotiated-price markets," *Journal of Political Economy*, 2019, *127* (4), 1550–1598.

Armstrong, Mark, "Search and ripoff externalities," *Review of Industrial Organization*, 2015, *47* (3), 273–302.

Asker, John, "A study of the internal organization of a bidding cartel," *The American Economic Review*, 2010, *100* (3), 724–762.

Athey, Susan and Philip A Haile, "Identification of standard auction models," *Econometrica*, 2002, *70* (6), 2107–2140.

Bajari, Patrick, "Comparing competition and collusion: a numerical approach," *Economic Theory*, 2001, *18* (1), 187–205.

_ **and Gregory Lewis**, "Procurement contracting with time incentives: theory and evidence," Technical Report, National Bureau of Economic Research 2009.

Burdett, Kenneth and Kenneth L Judd, "Equilibrium price dispersion," *Econometrica: Journal of the Econometric Society*, 1983, pp. 955–969.

Commission, Citizen Budget, "Taxes In, Garbage Out The Need for Better Solid Waste Disposal Policies in New York City," 2012.

Efron, Bradley and Robert J Tibshirani, *An introduction to the bootstrap*, CRC press, 1994.

Gavazza, Alessandro, "An empirical equilibrium model of a decentralized asset market," *Econometrica*, 2016, *84* (5), 1755–1798.

Gehrig, Thomas, "Intermediation in search markets," *Journal of Economics & Management Strategy*, 1993, *2* (1), 97–120.

Goldberg, Pinelopi Koujianou, "Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey," *Journal of Political Economy*, 1996, pp. 622–654.

Gonzalez, José Luis Moraga, Zsolt Sándor, and Matthijs R Wildenbeest, "Prices, Product Differentiation, and Heterogeneous Search Costs," 2014.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong, "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 2000, *68* (3), 525–574.

Haile, Philip A, "Auctions with resale markets: an application to US forest service timber sales," *American Economic Review*, 2001, pp. 399–427.

_, Han Hong, and Matthew Shum, "Nonparametric tests for common values at first-price sealed-bid auctions," Technical Report, National Bureau of Economic Research 2003.

Hendel, Igal, Aviv Nevo, and François Ortalo-Magné, "The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison. com," *American Economic Review*, 2009, *99* (5), 1878–1898.

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

Honka, Elisabeth, "Quantifying search and switching costs in the US auto insurance industry," *The Rand Journal of Economics*, 2014, *45* (4), 847–884.

— **and Pradeep K Chintagunta**, "Simultaneous or sequential? search strategies in the us auto insurance industry," *Search Strategies in the US Auto Insurance Industry (April 29, 2014)*, 2014.

Hortaçsu, Ali and Chad Syverson, "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*, 2004, *119* (2), 403–456.

Jacobs, James B, Coleen Friel, and Robert Raddick, *Gotham unbound: How New York city was liberated from the grip of organized crime*, NYU Press, 2001.

Jin, Ginger Zhe and Phillip Leslie, "THE EFFECT OF INFORMATION ON PRODUCT QUALITY: EVIDENCE FROM RESTAURANT HYGIENE GRADE CARDS.," *Quarterly Journal of Economics*, 2003, *118* (2).

Kelly, Robert J, "Dirty Dollars: The Carting and Waste Disposal Industries," in "The Upperworld and the Underworld," Springer, 1999, pp. 135–155.

Krasnokutskaya, Elena, "Identification and estimation of auction models with unobserved heterogeneity," *The Review of Economic Studies*, 2011, *78* (1), 293–327.

— **and Katja Seim**, "Bid preference programs and participation in highway procurement auctions," *The American Economic Review*, 2011, *101* (6), 2653–2686.

Krishna, Vijay, *Auction theory*, Academic press, 2009.

Li, Tong and Quang Vuong, "Nonparametric estimation of the measurement error model using multiple indicators," *Journal of Multivariate Analysis*, 1998, *65* (2), 139–165.

— **and Xiaoyong Zheng**, "Entry and competition effects in first-price auctions: theory and evidence from procurement auctions," *The Review of Economic Studies*, 2009, *76* (4), 1397–1429.

Lizzeri, Alessandro, "Information revelation and certification intermediaries," *The RAND Journal of Economics*, 1999, pp. 214–231.

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

Matzkin, Rosa L, "Nonparametric identification," *Handbook of Econometrics*, 2007, *6*, 5307–5368.

McCall, John Joseph, "Economics of information and job search," *The Quarterly Journal of Economics*, 1970, pp. 113–126.

Meilijson, Isaac, "Estimation of the lifetime distribution of the parts from the autopsy statistics of the machine," *Journal of Applied Probability*, 1981, pp. 829–838.

Nguyen, Thuy TT and Bruce G Wilson, "Fuel consumption estimation for kerbside municipal solid waste (MSW) collection activities," *Waste Management & Research*, 2010, *28* (4), 289–297.

Rubinstein, Ariel and Asher Wolinsky, "Middlemen," *The Quarterly Journal of Economics*, 1987, pp. 581–594.

Rust, John and George Hall, "Middlemen versus Market Makers: A Theory of Competitive Exchange," *Journal of Political Economy*, 2003, *111* (2).

Salop, Steven and Joseph Stiglitz, "Bargains and ripoffs: A model of monopolistically competitive price dispersion," *The Review of Economic Studies*, 1977, pp. 493–510.

Sorensen, Alan T, "Equilibrium price dispersion in retail markets for prescription drugs," *Journal of Political Economy*, 2000, *108* (4), 833–850.

Spulber, Daniel F, "Market making by price-setting firms," *The Review of Economic Studies*, 1996, *63* (4), 559–580.

—, "Market microstructure and intermediation," *The Journal of Economic Perspectives*, 1996, pp. 135–152.

Stahl, Dale O, "Oligopolistic pricing with sequential consumer search," *The American Economic Review*, 1989, pp. 700–712.

Stigler, George J, "The economics of information," *The journal of political economy*, 1961, pp. 213–225.

A Data Coding

A.1 Sample Construction

A.2 Definition of Brokered Contract and Number of Bidders

This section lays out the details of data coding and sample selection.

- *Definition of Brokered Contracts:* There are two data fields from which a brokered contract can be inferred. The first one is a field that contains an indicator on how the contract was formed. A second field provides the name of the brokerage company. A contract is defined as brokered if either the first field indicates so or if the latter field contains an entry that does not disqualify as a broker company. A contract is defined as brokered if the broker indicator is one in either the first or second reporting period. Because I am only using the initial contract period I am not using any information broker variable after the first year. The broker firm identifier is hand-coded from the second field, which is a string variable. In most cases it is straightforward to assign a given observation to a recurring broker-name. In some cases, however, the listed name might be an individual contact name and therefore not allow me to match this observation to a broker company. In this case, I would infer that such a "broker" has only one contract with one carter. Since the broker identifier is used to infer the number of bidders, I proceed with observations that have just one carter listed as follows: if a broker name only appears once I randomly impute the number of bidders from the remaining distribution of observations that can be unambiguously assigned. The data also shows whether a contract is closed after solicitation through a carter. I am currently ignoring this information although it could be interesting for future work to separately identify the search cost from haggling cost.

- *Sample restrictions:* The dataset has a panel structure. Each customer-carter relationship is documented biannually, leading to 1,184,641 panel observations. Out of those I drop contracts with missing information on carters, prices, or quantity, which eliminates 256,874 observations. I only use contracts whose prices are quoted in volume. Discarding weights based contracts drops 237,334 observations. Due to its isolated geographic position I further drop contracts from Staten Island, which amounts to 42,692 observations. I drop contracts with missing quantity information, which drops 256,738 observations. After these restrictions I am left with 941,770 panel observations. I then only keep the first reporting period for new contracts, which leaves me with 170,232 initial relationships. Out of those I only keep contracts where the buyer falls within the retail category, leaving me with a final sample of 44417 observations.

- *Winsorizing:* I winsorize the rate variable at the 1st and 99th percentile to deal with outliers. I winsorize the quantity variable at the 1st and 99th percentile in all reduced form exercises. For the structural exercise I further winsorize the quantity variable at the 95th percentile to prevent the average in the highest quantity bin to be driven by a few very large customers.

B Additional Details on the Identification

B.1 Illustrating lack of identification from prices alone

Figure 7 illustrates why prices in the search-market alone would not be enough to tell search cost and carters' service cost apart. It shows simulated data from two markets. The distribution of contract prices, shown in the top left corner, are the same, and the two markets are therefore observationally equivalent. The underlying primitives that give rise to these distributions are, however, very different. In

Figure (7) Observationally equivalent markets with different primitives.

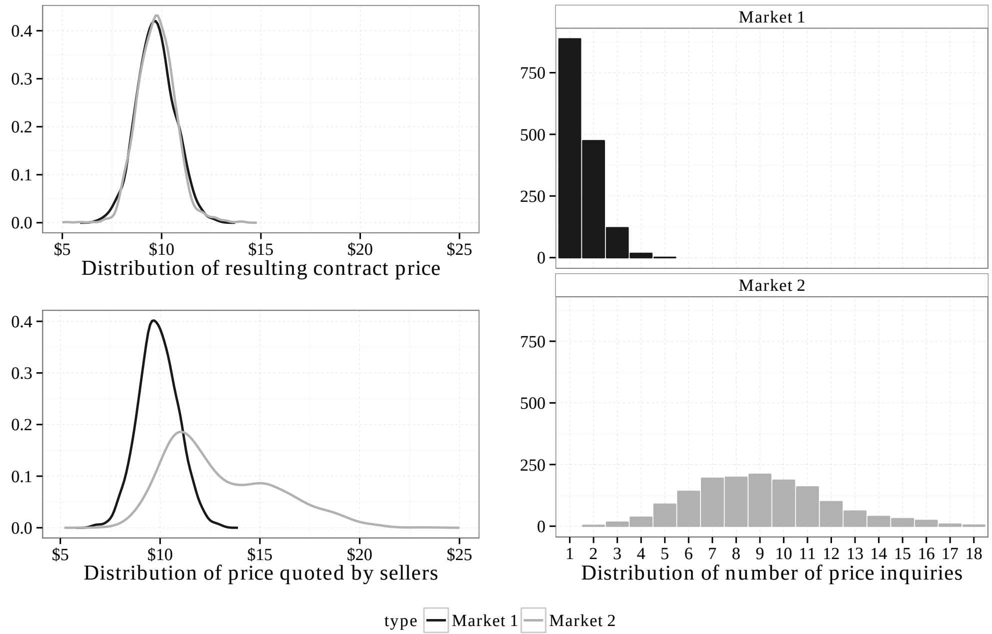

Notes: The four panels show data from two *simulated* markets. The top left panel is the contract price distribution observed by the econometrician and the bottom left panel the price offer distributions. The two right panels show the distribution of price inquiries for the two markets.

market 1 buyers have higher search cost and ask for fewer price quotes relative to market 2. Market 1, on the other hand, is served by firms that have lower cost and quote lower prices. This example shows that it will not be enough to observe just one price distribution to tell apart the search cost for buyers and the cost of service provision for carters.

B.2 Identification of Consumer Search Cost

Proof of **Proposition 1:** The dependence of the distribution functions for the unobserved objects on observables is omitted for the proof.

Step 1: Get Carter Cost from Broker Market

Building on Guerre et al. (2000) we can recover the cost of carters based on the mapping between the distribution of observed winning bids $\mathcal{F}_b^o(\cdot)$ and cost. For the remaining argument we can therefore treat the carter cost distribution $\mathcal{G}(\cdot)$ as known.

$$c = p^b - \frac{1 - \mathcal{F}_b^o\left(p^b|N_b\right)}{(N_b - 1) \cdot f_b^o\left(p^b|N_b\right)} \tag{11}$$

B.2.1 Step 2: Show that the Bidding Function is Identified given $\mathcal{G}(\cdot)$

The second step of the argument shows that given the observed prices and the cost distribution, one can recover the bidding function uniquely. The bidding function is as following

$$\beta_S(c) = \arg\max_p (p - c) \sum_m w_m \left[1 - \mathcal{G}(\beta_S^{-1}(p))\right]^{m-1}. \tag{12}$$

and the distribution of prices is observable, with

$$1 - \mathcal{F}^o(p) = \sum_m w_m \left[1 - \mathcal{G}(\beta_S^{-1}(p))\right]^m. \tag{13}$$

The first order condition of the bidding problem in (12) is

$$\begin{aligned}&\sum_m w_m[1 - \mathcal{G}(\beta_S^{-1}(p))]^{m-1} \\ &+(p - c)\sum_m w_m(m-1)[1 - G(\beta_S^{-1}(p))]^{m-2}\left[-\frac{d}{dp}\mathcal{G}(\beta_S^{-1}(p))\right] = 0.\end{aligned} \tag{14}$$

Define

$$\mathcal{A}(p) = \sum_m w_m[1 - \mathcal{G}(\beta_S^{-1}(p))]^{m-1}. \tag{15}$$

From Equation (13),

$$\mathcal{A}(p) = \frac{1 - \mathcal{F}^o(p)}{1 - \mathcal{G}(\beta_S^{-1}(p))}, \tag{16}$$

and from Equation (14),

$$\mathcal{A}(p) = (p - \beta_S^{-1}(p)) \sum_m w_m (m-1)[1 - \mathcal{G}(\beta_S^{-1}(p))]^{m-2} \left[\frac{d}{dp} \mathcal{G}(\beta_S^{-1}(p)) \right]. \quad (17)$$

Furthermore, taking the derivative of (13) with respect to p and decomposing the resulting term give

$$f^o(p) = \sum_m w_m (m-1)[1 - \mathcal{G}(\beta_S^{-1}(p))]^{m-1} \left[\frac{d}{dp} \mathcal{G}(\beta_S^{-1}(p)) \right] + \sum_m w_m [1 - \mathcal{G}(\beta_S^{-1}(p))]^{m-1} \left[\frac{d}{dp} \mathcal{G}(\beta_S^{-1}(p)) \right].$$

Using Equation (17) for the first term and Equation (16) for the second term, we can rewrite the above equation as

$$f^o(p) = \frac{\mathcal{A}(p)[1 - \mathcal{G}(\beta_S^{-1}(p))]}{p - \beta_S^{-1}(p)} + \mathcal{A}(p) \left[\frac{d}{dp} \mathcal{G}(\beta_S^{-1}(p)) \right].$$

Using Equation (16) we can further simplify the above equation to

$$f^o(p) = \frac{1 - \mathcal{F}^o(p)}{p - \beta_S^{-1}(p)} + [1 - \mathcal{F}^o(p)] \frac{g(\beta_S^{-1}(p))}{1 - \mathcal{G}(\beta_S^{-1}(p))} \left[\frac{d}{dp} \beta_S^{-1}(p) \right].$$

Dividing by $1 - \mathcal{F}^o(p)$ gives

$$\frac{f^o(p)}{1 - \mathcal{F}^o(p)} = \frac{g(\beta_S^{-1}(p))}{1 - \mathcal{G}(\beta_S^{-1}(p))} \left[\frac{d}{dp} \beta_S^{-1}(p) \right] + \frac{1}{p - \beta_S^{-1}(p)}.$$

Define the inverse bidding function as $h(p)$ then the above equality defines the following ODE

$$h'(p) = \left[\frac{f^o(p)}{1 - \mathcal{F}^o(p)} - \frac{1}{p - h(p)} \right] \frac{1 - \mathcal{G}(h(p))}{g(h(p))} =: \phi(p, h(p)).$$

The initial condition is $h(\underline{p}) = \underline{c}$, where $\underline{p}$ is the lowest price in the support of $\mathcal{F}^o(p)$ and $\underline{c}$ is the lowest cost in the support of $\mathcal{G}(c)$. Since we only need $\beta_S(p)$ on some interval to identify the weights, the ODE only needs to be unique on some interval. To argue formally for the uniqueness of $h(p)$ on some interval $[\underline{p}, \tilde{p}]$ for some $\tilde{p} > \underline{p}$, pick some p^* and c^{**} such that $c < c^{**} < p < p^*$ and $\mathcal{F}^o(p^*) < 1$. By the Picard-Lindeløf Theorem the sufficient conditions for existence and uniqueness are that (i)

$\phi(p, h)$ is continuous in $R = \left[\underline{p}, p^*\right] \times [\underline{c}, c^{**}]$, and (ii) ϕ is Lipschitz-continuous in h, that is for some $L > 0$: $|\phi(p, h_1) - \phi(p, h_2)| \leq L|h_1 - h_2| \quad \forall (p, h_1), (p, h_2) \in R$. Let K be an upper bound of $\phi(p, h)$ in R, then the ODE has a unique solution $h = h(p)$ defined on:

$$\left[\underline{p}, \underline{p} + \frac{c^{**} - \underline{c}}{K}\right] =: [\underline{p}, \tilde{p}]$$

Being able to identify $h(p)$ on a non-degenerate interval is enough for our purpose. Since this function is identified on $[\underline{p}, \tilde{p}]$ and it should be strictly increasing if the model is correct, the bidding function $\beta = h^{-1}$ is identified on $[\underline{c}, \tilde{c}]$ for $\tilde{c} = h(\tilde{p})$. The argument for identifying the weights from the bidding function still follows when I use the bidding function restricted to a proper interval. Lastly, I will show that Lipschitz-continuity holds under some elementary condition, which is that:

$$\frac{1 - \mathcal{G}(h)}{g(h)}$$

is Lipschitz-continuous on $[c, c^{**}]$ with parameter L_1. Suppose that this holds true. Fix any $p \in \left[\underline{p}, p^*\right]$. For any $h_1, h_2 \in [\underline{c}, c^{**}]$, then:

$$\begin{aligned}\phi(p, h_1) - \phi(p, h_2) &= \frac{f^o(p)}{1 - \mathcal{F}^o(p)} \left[\frac{1 - \mathcal{G}(h_1)}{g(h_1)} - \frac{1 - \mathcal{G}(h_2)}{g(h_2)}\right] \\ &\quad - \frac{1}{(p - h_1)(p - h_2)} \left[(p - h_2) \frac{1 - \mathcal{G}(h_1)}{g(h_1)} - (p - h_1) \frac{1 - \mathcal{G}(h_2)}{g(h_2)}\right] =\end{aligned}$$

$$\begin{aligned}&\frac{f^o(p)}{1 - \mathcal{F}^o(p)} \left[\frac{1 - \mathcal{G}(h_1)}{g(h_1)} - \frac{1 - \mathcal{G}(h_2)}{g(h_2)}\right] - \frac{p}{(p - h_1)(p - h_2)} \left[\frac{1 - \mathcal{G}(h_1)}{g(h_1)} - \frac{1 - \mathcal{G}(h_2)}{g(h_2)}\right] \\ &\quad - \frac{1}{(p - h_1)(p - h_2)} \left[(h_1 - h_2) \frac{1 - \mathcal{G}(h_1)}{g(h_1)} - h_1 \left(\frac{1 - \mathcal{G}(h_1)}{g(h_1)} - \frac{1 - \mathcal{G}(h_2)}{g(h_2)}\right)\right]\end{aligned}$$

Therefore:

$$|\phi(p,h_1)-\phi(p,h_2)| \le \Bigg\{ \left(\max_{p\in[\underline{p},p^*]} \frac{f^o(p)}{1-F(p)} \right) L_1 +$$

$$\frac{1}{\left(\underline{p}-c^{**}\right)^2} \left[p^* L_1 + \left(\max_{h\in[c,c^{**}]} \frac{1-\mathcal{G}(h)}{g(h)} \right) + c^{**} L_1 \right] \Bigg\} |h_1-h_2|$$

Defining the terms in curly brackets as L completes this part of the proof.

B.2.2 Step 3: Identify Weights from Bidding Function $\beta_S()$

Suppose the cost function $\mathcal{G}(\cdot)$ and the bidding function $\beta_S(\cdot)$ are known. Suppose that two sets of weights $\{w_m\}_{m=1}^M$ and $\{\lambda_m\}_{m=1}^M$ both give rise to the bidding function. This means for all c in the support:

$$\frac{\sum_m w_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}}{\sum_m w_m (1-\mathcal{G}(c))^{m-1}} = \beta_S(c) - c = \frac{\sum_m \lambda_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}}{\sum_m \lambda_m (1-\mathcal{G}(c))^{m-1}}. \tag{18}$$

Multiplying the terms leads to

$$\begin{aligned} LHS(c) &= \left[\sum_m w_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}\right] \left[\sum_m \lambda_m (1-\mathcal{G}(c))^{m-1}\right] \\ &= \left[\sum_m \lambda_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}\right] \left[\sum_m w_m (1-\mathcal{G}(c))^{m-1}\right] = RHS(c) \end{aligned} \tag{19}$$

Taking derivatives with respect to cost,

$$\begin{aligned} \frac{d}{dc} LHS(c) &= -\left[\sum_m w_m (1-\mathcal{G}(c))^{m-1}\right] \left[\sum_m \lambda_m (1-\mathcal{G}(c))^{m-1}\right] \\ &\quad - \left[\sum_m w_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}\right] \left[\sum_m (m-1)\lambda_m (1-\mathcal{G}(c))^{m-2} g(c)\right] \end{aligned}$$

and

$$\begin{aligned} \frac{d}{dc} RHS(c) &= -\left[\sum_m \lambda_m (1-\mathcal{G}(c))^{m-1}\right] \left[\sum_m w_m (1-\mathcal{G}(c))^{m-1}\right] \\ &\quad - \left[\sum_m \lambda_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}\right] \left[\sum_m (m-1) w_m (1-\mathcal{G}(c))^{m-2} g(c)\right]. \end{aligned}$$

Combining the derivatives gives

$$\frac{\sum_m (m-1) w_m (1-\mathcal{G}(c))^{m-2}}{\sum_m (m-1)\lambda_m (1-\mathcal{G}(c))^{m-2}} = \frac{\sum_m w_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}}{\sum_m \lambda_m \int_c^{\bar{c}} (1-\mathcal{G}(\tilde{c}))^{m-1} d\tilde{c}} = \frac{\sum_m w_m (1-\mathcal{G}(c))^{m-1}}{\sum_m \lambda_m (1-\mathcal{G}(c))^{m-1}},$$

where the last equality follows from Equation (18). This can be simplified to

$$\sum_m \sum_n (m-1) w_m \lambda_n (1-\mathcal{G}(c))^{m+n} = \sum_m \sum_n (n-1) w_m \lambda_n (1-\mathcal{G}(c))^{m+n}. \tag{20}$$

Note that each side of the above equation is a polynomial in $1-\mathcal{G}(c)$. If this equation holds on some interval of c and thus some interval of $1-\mathcal{G}(c)$ (given that $\mathcal{G}(\cdot)$ is strictly increasing so that the interval is not degenerate), then by the unisolvence theorem for polynomials, the coefficients of each power term must equate.

With $m+n=3$,

$$w_2\lambda_1 = w_1\lambda_2 \iff \frac{w_2}{\lambda_2} = \frac{w_1}{\lambda_1} =: r.$$

Suppose that it holds for all $l < L \leq M$ that $\frac{w_l}{\lambda_l} = r$. Equating the coefficients for $m+n = L+1$, we have

$$(L-1) w_L \lambda_1 + \sum_{m=2}^{L-1} \sum_{n=1}^{L+1-m} (m-1) w_m \lambda_n = (L-1) w_1 \lambda_L + \sum_{n=2}^{L-1} \sum_{m=1}^{L+1-n} (n-1) w_m \lambda_n.$$

Hence:

$$\frac{w_L}{\lambda_L} = \frac{w_1}{\lambda_1} = r$$

That is, it holds inductively that $\frac{w_n}{\lambda_n} = \frac{w_m}{\lambda_m}$ for all m, n. This, together with $\sum_m w_m = \sum_m \lambda_m = 1$ means that $w_m = \lambda_m$ for all m. This concludes the proof that the weights are uniquely identified.

B.3 Dealing with Asymmetries

In this section I briefly describe a setup that could accommodate persistent cost differences across carters. Such a setup would be important if one was worried about persistent cost differences across carters in the broker market and in the search market. However, because the bidding functions in this setup do not have a closed form solution, the identification strategy of the search-weights does not extent to this setup. This selection will be accounted for by allowing for different cost distributions, which are going to be indexed by L and H. The classification into H and L would have to be based on observables, for instance the density of the network, carter size, etc. One could then adjust for differences in composition of high and low cost firms across the two markets. The maximization problem in the broker auction and in the search market is given by the following two equations:

$$\max_p (p-c) \cdot \widetilde{\mathcal{G}}_L(\beta_{b,L}^{-1}(p))^{N_{b,L}} \cdot \widetilde{\mathcal{G}}_H(\beta_{b,H}^{-1}(p))^{N_{b,H}-1}$$

and

$$\max_p (p-c) \cdot \left[\sum_{m=1}^{M-1} w_m \cdot \sum_{k=0}^{m} \cdot \frac{\binom{N_L}{k} \cdot \binom{N_H-1}{m-k}}{\binom{N_H+N_L-1}{m}} \cdot \widetilde{\mathcal{G}}_L(\beta_{S,L}^{-1}(p))^k \cdot \widetilde{\mathcal{G}}_H(\beta_{S,H}^{-1}(p))^{m-k} \right].$$

Figure 8 illustrates the type of selection that one could account for. Since the cost distributions are recovered from the brokered market, identification requires that in this part of the market there are infinitely many auctions with at least one firm of each type.

Carter costs are identified according to **Theorem 6** in Athey and Haile (2002), which states that the distributions of valuations in asymmetric auctions can be identified from the transaction price and the identity of the winner.[38] Therefore, one could treat $\mathcal{G}_k(\cdot|z), k \in \{L, H\}$ as known objects for all firms that bid in the brokered market. Because the identity of the firms is known and under the assumption of full support, i.e. that one observed infinitely many auctions for each type of firm in the broker market, one could simply take the identified cost functions from the broker market and adjust for the change in composition in L and

[38] This is an application of the identification result in competing risk models presented in Meilijson (1981)

Figure (8) Selection in The Market

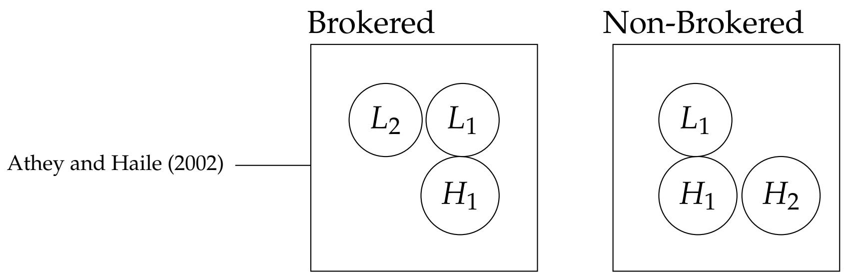

Notes: This figure provides a graphical illustration of the type of selection that is allowed on the seller side. Firms L and H are both observed in the brokered markets. Due to arguments presented by Athey and Haile (2002), the costs are identified from observed contract prices in the brokered market. This, in turn, allows us to assign the already known cost functions from the broker market to firms of type L and H in the search market.

H firms in the search market. However, because the resulting bidding functions are asymmetric, one would have to solve for equilibria using the approach in Bajari (2001), which would require an approximation of the bidding function and the first order condition through a polynomial. Bernstein polynomia are a natural candidate for such an approximation because they allow to impose shape constraints.

C Dynamic Development of Prices

This project focuses on the static aspects of pricing in this market. Because customers enter a relationship with carters they might be forward looking with regard to the prices that carters charge in the future and this might affect their initial decision to search. Sellers might, for example, engage in bait and switch pricing, luring customers with initial low rates and then increase their prices when the customer is "locked in". This strategy works if customers have significant switching cost. Another consideration might be that customers learn about the suppliers as time passes and decreasing search cost is reflected in a decreasing price path. To give a sense of the price developments of contracts across time I show results from a regression that includes the same controls as the price dispersion regressions:

$$\mathbf{X} = \{\text{business type FE, recyclables FE, time FE, zip code FE, transfer station FE,} \\ q, \ldots, q^5 \text{Number of Pickup FE}\}.$$

The variable of interest is now a time dummy, which summarizes how prices move across time holding the above characteristics fixed. These dummies show how prices develop within a customer carter relationship. Effects are shown relative to the excluded category, which is the initial price. The reporting periods are six month. We can see that there is a small dip in the second reporting period and then a sustained change in prices after two years, which is the time at which customers must be able to switch without any penalty according to regulation.

Figure (9) Dynamic development of prices

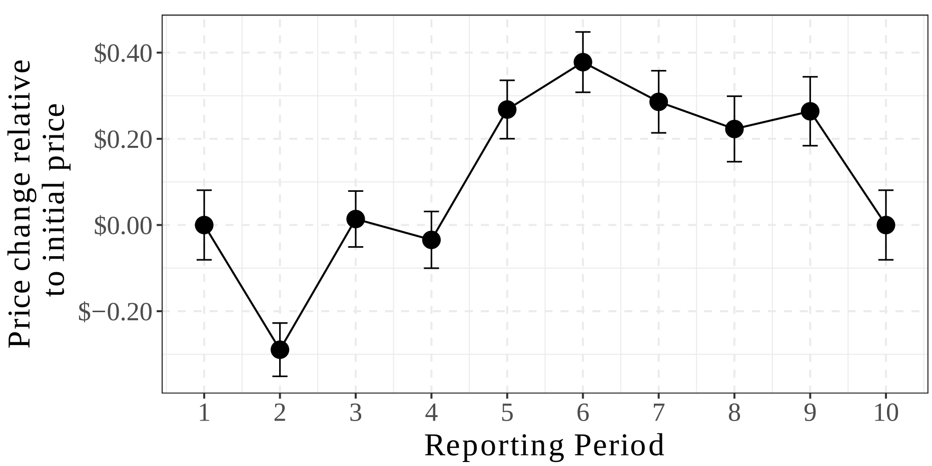

Notes: This figure shows how the prices develop for buyers over time, relative to the initial price. A reporting period is a half-year. The data points are time-dummy estimates, where the excluded variable is the dummy on the initial time period.

D Additional Graphs

D.1 Quantity Variation and Model Fit

The following graph illustrates why it would be hard to fit the variation across consumers with different quantity if the search cost can not depend directly on q. The graph shows the cut-off types both for buyers with $q_1 < q_2$. Buyers with q_2 have higher incentives to search and their cut-off types in terms of search cost therefore move up. Suppose that the true fraction of brokered contracts for buyers with q_2 is given by the gray area. Given a particular search cost distribution,

Figure (10) Restriction on Search Cost and Brokered Contracts

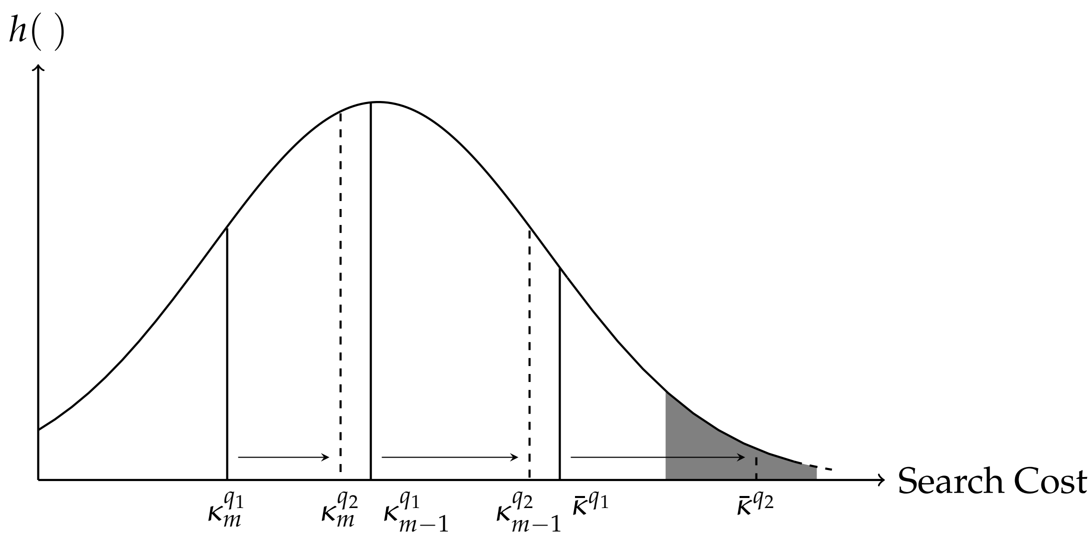

the predicted cut-off type would be given by $\bar{\kappa}^{q_2}$, which would under-predict the fraction of brokered contracts. This illustrates that without letting the search cost distribution directly depend on q it would be hard to fit the fraction of buyers that go to a broker under a given parametric assumption, which disciplines the tails of the distribution. I will therefore allow the selection margin to brokers depend on q. Letting the relative attractiveness of brokers depend on q has a natural interpretation in this context. Larger more important buyers might, for example, receive additional services. Note that the alternative would be to estimate a search cost distribution for each q. However, since the search cost has the interpretation of a marginal cost, this is the less plausible interpretation.

E Additional Tables

E.1 Breakdown of Business-types

Table (8) Types of Businesses signed up with carters

Business type	Fraction of Total	Total Number
Retail non - food	.410	46514
Retail - food	.138	15660
Wholesale non - food	.015	1648
Wholesale - food	.010	1048
Restaurant/bar	.107	12181
Hotel - small	.003	302
Hotel-big	.002	175
Medical offices	.028	3138
Automobile repair	.027	3100
Office building - small	.032	3620
Office building - medium	.026	2997
Office building - large	.015	1675
Light manufacturing	.020	2242
Heavy manufacturing	.003	298
Institution	.020	2277
Professional office	.043	4838
None of the above	.104	11883

Note: Fractions and totals are averaged across the nine reporting periods. Missings are declared as "none of the above".

Table (9) Fraction of Contracts awarded through Broker.

Variable	Brokered	Variable	Brokered
Types of business		**Borough**	
Food, Retail/Wholesale	14%	Bronx	17%
Restaurants	15%	Brooklyn	14%
Hotels	14%	Manhattan	14%
Office Building	9%	Queens	9%
Manufacturing/Repair	10%	**Quantity (percentile)**	
Institution (Hospital/University)	40%	$q < 25\%$	6%
Recyclables		$25\% < q < 50\%$	15%
No	13%	$50\% < q < 75\%$	14%
Yes	14%	$q > 75\%$	19%

Note: This table shows conditional percentages of broker usage.

Table (10) Documenting differences between brokered and un-brokered contracts. Dependent rate-variable does not include broker commissions.

	(1)	(2)	(3)	(4)	(5)
	OLS	OLS	Quantile 0.25	Quantile 0.5	Quantile 0.75
	p_{ijt} (rate)	p_{ijt} (rate)	p_{ijt} (rate)	p_{ijt} (rate)	p_{ijt} (rate)
Contract brokered	-1.550**	-1.559**	-1.775**	-1.000**	-0.0821
	(0.158)	(0.160)	(0.0863)	(0.0404)	(0.0541)
q	-0.0113**	-0.0114**	-0.000431	2.22e-16	-0.000437
	(0.00159)	(0.00159)	(0.000503)	(0.000236)	(0.000316)
Recyclables	0.451**	0.459**	0.520**	0.200**	0.0867*
	(0.153)	(0.151)	(0.0689)	(0.0322)	(0.0432)
Carter deals /w broker		0.148			
		(0.193)			
Observations	35659	35659	38135	38135	38135
Deals with Broker	No	Yes	No	No	No
Transfer FE	Yes	Yes	No	No	No
R^2	0.312	0.312			

Note: $+ \ p < 0.10$, $* \ p < 0.05$, $** \ p < 0.01$. All specifications include the following set of controls: quantity of waste, transfer-station fixed effects, zip-code fixed effects, business-type fixed effects, length-of-contract fixed effects, recyclable materials fixed effects, reporting-date fixed effects, number of weekly pickups, and the HHI index. The aggregate regressions at the zip-code level include the average quantity at the zip-code level, the average number of pickups and the average number of customers that use recyclables. Standard errors are clustered at the zip-code level.

Table (11) Documenting differences between brokered and un-brokered contracts. Dependent rate-variable includes broker commissions.

	(1)	(2)	(3)	(4)	(5)
	OLS	OLS	Quantile 0.25	Quantile 0.5	Quantile 0.75
	p_{ijt} (rate + fee)	p_{ijt} (rate + fee)	p_{ijt} (rate + fee)	p_{ijt} (rate + fee)	p_{ijt} (rate + fee)
Contract Brokered	1.912**	1.901**	0.242**	2.590**	4.732**
	(0.406)	(0.407)	(0.0882)	(0.0516)	(0.0620)
q	-0.0139**	-0.0139**	-0.00142**	2.44e-15	-0.000335
	(0.00207)	(0.00207)	(0.000514)	(0.000301)	(0.000361)
Recyclables	0.410*	0.420*	0.594**	0.180**	0.0724
	(0.177)	(0.177)	(0.0704)	(0.0411)	(0.0494)
Carter deals /w broker		0.185			
		(0.196)			
Observations	35659	35659	38135	38135	38135
Deals with Broker	No	Yes	No	No	No
Transfer FE	Yes	Yes	No	No	No
R^2	0.324	0.324			

Note: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$. All specifications include the following set of controls: quantity of waste, transfer-station fixed effects, zip-code fixed effects, business-type fixed effects, length-of-contract fixed effects, recyclable materials fixed effects, reporting-date fixed effects, number of weekly pickups, and the HHI index. The aggregate regressions at the zip-code level include the average quantity at the zip-code level, the average number of pickups and the average number of customers that use recyclables. Standard errors are clustered at the zip-code level.

F Restrictions Plot

These two plot show all the restrictions that the recovered weights impose on the search cost distribution. Each line in the graph corresponds has a corresponding known empirical CDF that is derived from the recovered weights. The probability mass between each consecutive lines is therefore known. The firs graph shows how finely the search cost distribution is identified without using variation in q. The second graph shows all the marginal types across buyers with different quantities. Since the quantity is not directly affecting the search cost, variation in q imposes even more restrictions.

Figure (11) Restrictions Plot

(a) Marginal Types holding Quantity Fixed

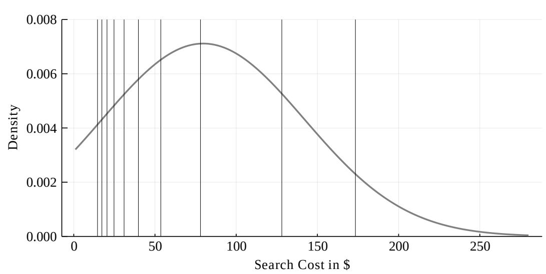

(b) Exploiting all Marginal Types

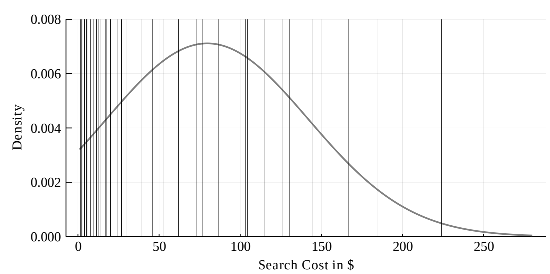

Figure (12) Overview of Changes when Brokers are not available

Subset	Price Change (%)	Baseline Brokered Expenses ($)	Baseline Not Brokered Expenses ($)	Carter Profits (QW) ($)
Bronx				
$Q_{0,25}$	0.25	28.95	0.29	-0.28
$Q_{25,50}$	0.27	175.13	0.19	2.4
$Q_{50,75}$	6.46	267.53	123.89	29.31
$Q_{75,100}$	4.32	1406.77	138.53	-74.13
Brooklyn				
$Q_{0,25}$	0.19	177.85	0.16	0.1
$Q_{25,50}$	0.71	319.46	2.25	3.72
$Q_{50,75}$	18.3	1220.0	164.4	66.76
$Q_{75,100}$	6.07	5217.15	128.87	-145.34
Manhattan				
$Q_{0,25}$	0.22	50.47	0.13	-0.13
$Q_{25,50}$	0.29	95.82	-0.14	1.56
$Q_{50,75}$	5.57	198.9	110.23	23.05
$Q_{75,100}$	3.49	1295.03	97.19	-55.44
Queens				
$Q_{0,25}$	0.19	419.11	0.16	0.1
$Q_{25,50}$	0.62	652.6	1.08	4.12
$Q_{50,75}$	18.28	2431.05	139.02	66.85
$Q_{75,100}$	6.09	3729.27	140.45	-145.85

Note: This table gives an overview over absolute and percentage changes in the total cost (search cost + contract cost). For buyers that were formerly already transacting directly with carters, the table provides a further breakdown into the changes in search expenses and contract expenses. For buyers in the broker market such a breakdown does not make much sense since the initial broker expenses provide no logical separation into search and contract expenses. All numbers are computed over the length of a contract, which is two years.

G Equilibrium Uniqueness

In this section I show some examples to demonstrate that multiple equilibria do not appear to be an issue in this setting. In each case, the first column shows the starting set of weights and I show the maximal number of iterations that are needed to satisfy the stopping criterion. The parameters that I use are the estimates that I recover but I could have used other parameters as well. What the examples below show is that for a wide variety of starting values, the process converges to the same set of weights and does so quickly.

Table (12) Example 1

Weights	It 1	It 2	It 3	It 4	It 5	It 6	It 7
w_1	0.0	0.286	0.219	0.201	0.198	0.197	0.197
w_2	1.0	0.266	0.258	0.257	0.257	0.256	0.256
w_3	0.0	0.148	0.16	0.163	0.164	0.164	0.164
w_4	0.0	0.086	0.098	0.101	0.101	0.101	0.102
w_5	0.0	0.051	0.06	0.062	0.063	0.063	0.063
w_6	0.0	0.037	0.044	0.045	0.046	0.046	0.046
w_7	0.0	0.022	0.027	0.028	0.028	0.029	0.029
w_8	0.0	0.021	0.025	0.026	0.027	0.027	0.027
w_9	0.0	0.013	0.017	0.017	0.018	0.018	0.018
w_{10}	0.0	0.07	0.094	0.098	0.099	0.1	0.1

Table (13) Example 2

Weights	It 1	It 2	It 3	It 4	It 5	It 6
w_1	0.0	0.24	0.208	0.199	0.197	0.197
w_2	0.0	0.267	0.258	0.257	0.257	0.256
w_3	1.0	0.158	0.162	0.164	0.164	0.164
w_4	0.0	0.094	0.099	0.101	0.101	0.101
w_5	0.0	0.057	0.061	0.063	0.063	0.063
w_6	0.0	0.04	0.045	0.046	0.046	0.046
w_7	0.0	0.025	0.028	0.028	0.029	0.029
w_8	0.0	0.023	0.026	0.027	0.027	0.027
w_9	0.0	0.015	0.017	0.017	0.018	0.018
w_{10}	0.0	0.082	0.096	0.099	0.1	0.1

Table (14) Example 3

Weights	It 1	It 2	It 3	It 4	It 5
w_1	0.0	0.195	0.197	0.197	0.197
w_2	0.0	0.261	0.257	0.257	0.256
w_3	0.0	0.166	0.164	0.164	0.164
w_4	0.0	0.102	0.101	0.101	0.101
w_5	1.0	0.063	0.063	0.063	0.063
w_6	0.0	0.045	0.046	0.046	0.046
w_7	0.0	0.029	0.028	0.029	0.029
w_8	0.0	0.026	0.027	0.027	0.027
w_9	0.0	0.017	0.018	0.018	0.018
w_{10}	0.0	0.097	0.099	0.1	0.1

Table (15) Example 4

Weights	It 1	It 2	It 3	It 4	It 5	It 6
w_1	0.0	0.166	0.19	0.195	0.196	0.196
w_2	0.0	0.253	0.256	0.256	0.256	0.256
w_3	0.0	0.169	0.166	0.164	0.164	0.164
w_4	0.0	0.107	0.103	0.102	0.102	0.102
w_5	0.0	0.068	0.064	0.063	0.063	0.063
w_6	0.0	0.049	0.046	0.046	0.046	0.046
w_7	0.0	0.032	0.029	0.029	0.029	0.029
w_8	1.0	0.028	0.027	0.027	0.027	0.027
w_9	0.0	0.019	0.018	0.018	0.018	0.018
w_{10}	0.0	0.11	0.101	0.1	0.1	0.1

Table (16) Example 5

Weights	It 1	It 2	It 3	It 4	It 5	It 6
w_1	0.0	0.182	0.187	0.194	0.196	0.196
w_2	0.0	0.222	0.255	0.256	0.256	0.256
w_3	0.0	0.169	0.166	0.165	0.164	0.164
w_4	0.0	0.11	0.103	0.102	0.102	0.102
w_5	0.0	0.069	0.064	0.063	0.063	0.063
w_6	0.0	0.05	0.047	0.046	0.046	0.046
w_7	0.0	0.033	0.029	0.029	0.029	0.029
w_8	0.0	0.029	0.027	0.027	0.027	0.027
w_9	0.0	0.02	0.018	0.018	0.018	0.018
w_{10}	1.0	0.116	0.103	0.1	0.1	0.1

Table (17) Example 6

Weights	It 1	It 2	It 3	It 4	It 5	It 6
w_1	0.113	0.146	0.183	0.194	0.196	0.196
w_2	0.175	0.239	0.255	0.256	0.256	0.256
w_3	0.026	0.171	0.167	0.165	0.164	0.164
w_4	0.101	0.112	0.104	0.102	0.102	0.102
w_5	0.032	0.072	0.065	0.063	0.063	0.063
w_6	0.082	0.053	0.047	0.046	0.046	0.046
w_7	0.166	0.034	0.03	0.029	0.029	0.029
w_8	0.157	0.031	0.027	0.027	0.027	0.027
w_9	0.105	0.021	0.018	0.018	0.018	0.018
w_{10}	0.043	0.121	0.103	0.1	0.1	0.1

H Robustness Number of Bidders

Table (18) Robustness Check: Number of Bidders

	Change in Buyer Expenses			Carter Margin			
	Not Brokered	Brokered	All	(Not Quantity Weighted)	Carter Profits	Total Welfare Lower Bound	Total Welfare Upper Bound
Baseline							
Δ **Absolute**	\$64.0	\$445.0	\$127.0	\$0.046	\$-11.1	\$4.28	\$12.61
Δ **Percent**	2.52%	11.7%	4.6%	1.95%	-1.8%	4.41%	14.22%
30% *More Bidders*							
Δ **Absolute**	\$67.0	\$374.0	\$118.0	\$0.0	\$26.25	\$8.1	\$15.6
Δ **Percent**	2.7%	11.57%	4.49%	5.11%	4.52%	9.33%	19.66%
30% *Fewer Bidders*							
Δ **Absolute**	\$72.0	\$407.0	\$127.0	\$0.0	\$-30.92	\$4.16	\$11.93
Δ **Percent**	2.93%	12.12%	4.86%	0.08%	-4.85%	4.58%	14.37%

Note: This table shows expected search cost per inquiry, the number of inquiries, as well as total expenses for search. Service costs are estimated to be negative for a few subsets with large quantities, and percentage changes, therefore, not well defined. Search cost changes are computed under the assumption that brokers' total variable profits are equal to their fixed cost, which provides a lower bound on the change. Bootstrapped (standard errors) and [confidence intervals] based on 200 iterations.

I Welfare Calculations

The total change introduced by the in-availability of brokers can be composed into different parts. Let CE_t^s be customers' cost in when they were buying in the search market initially. This equals search cost plus contract cost, $CE_t^s = SE_t^s + CC_t^s = SE_t^s + q \cdot p_t^s$. Likewise, let $CE_t^b = SE_t^b + CC_t^b = SE_t^b + q \cdot p_t^b - \psi \cdot q$ be the expenses in t for customers that were buying in the brokered market initially. The last term accounts for the relative attractiveness of brokers as a function of q, which I interpret as non-pecuniary benefits. Search cost in the broker market are defined as the fraction of expenses that are marked as fees, i.e. $SE_t^b = \phi \cdot p_t^b \cdot q$. Let firms realized profits for customers that were originally in the search market be $\pi_t^s = q \cdot (p_t^s - c_t^s)$ and $\pi_t^b = q \cdot (p_t^b - c_t^b)$ for those originally in the brokered market. Finally, let η be the share of customers who contract in the search market. The upper bound on the total welfare change, ΔUBW, can then be computed as:

$$\Delta UBW = \eta \cdot ((q \cdot (p_1^s - c_1^s) - SE_1^s - q \cdot p_1^s) - (q \cdot (p_0^s - c_0^s) - SE_0^s - q \cdot p_0^s)\Big]$$
$$+ (1-\eta) \cdot ((q \cdot (p_1^b - c_1^b) - SE_1^b - q \cdot p_1^b + \psi \cdot q) - (q \cdot (p_0^b - c_0^b) - SE_0^b - q \cdot p_0^b) - SE_0^b)$$

. In the last term an extra SE_0^b is subtracted since search cost in the brokered market are just a transfer from buyers to brokers. Prices are eliminated from this expression to obtain the final expression for the upper bound on the welfare change:

$$\Delta UBW = \eta \cdot (q \cdot c_0^s + SE_0^s - q \cdot c_1^s - SE_1^s) + (1-\eta) \cdot (q \cdot c_0^b - \psi \cdot q - q \cdot c_1^b - SE_1^b)$$

By subtracting the broker variable profits the lower bound on the changes is obtained:

$$\Delta LBW = \Delta UBW - (1-\eta) \cdot q \cdot \phi \cdot p_0^b$$

Since every buyer needs to contract in this market, total welfare in the market can be measured in the total cost to provide the service plus the search cost that the market incurs in order to produce matches between buyers and sellers.

J Fit Plots

The plots below show the fit of simulated moments versus data moments. Each graph shows the full set of five moments that are targeted in the estimation, conditional on a quantity bin. Overall, the model matched the moments well. However, for some bins it underestimates the broker price.

Figure (13) Model Fit Plot

(a) Customer quantity quartile: $Q_{0,0.25}$

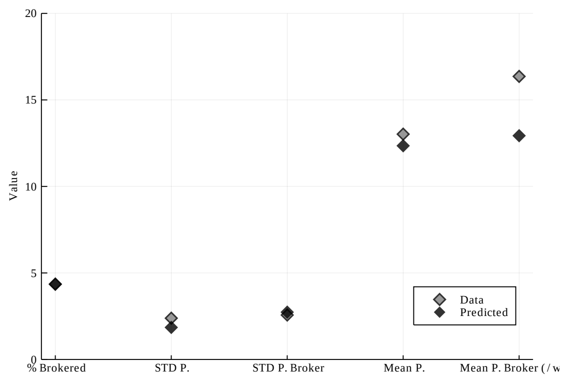

(b) Customer quantity quartile: $Q_{0.25,0.5}$

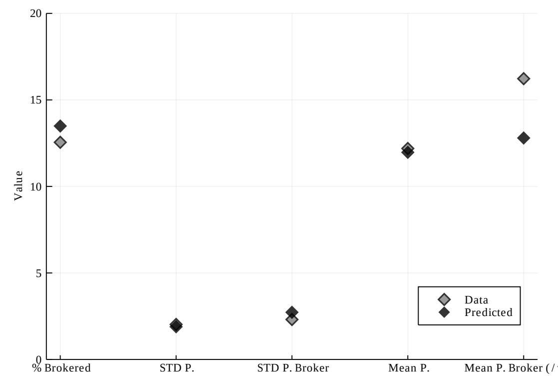

(c) Customer quantity quartile: $Q_{0.5,0.75}$

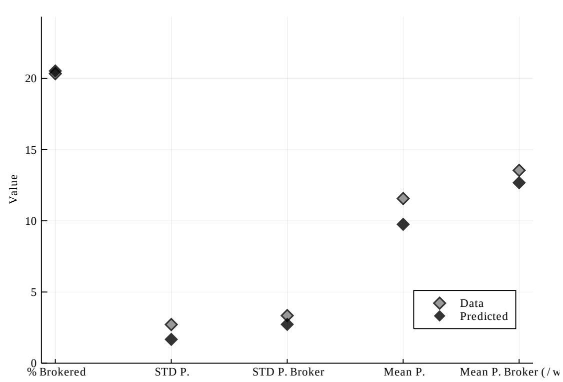

(d) Customer quantity quartile: $Q_{0.75,1.0}$