

MIT Open Access Articles

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

The MIT Faculty has made this article openly available. ***Please share***
how this access benefits you. Your story matters.

Citation: Salz, Tobias. 2022. "Intermediation and Competition in Search Markets: An Empirical Case Study." *The Journal of Political Economy*, 130 (2).

As Published: 10.1086/717349

Publisher: University of Chicago Press

Persistent URL: <https://hdl.handle.net/1721.1/145201>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Intermediation and Competition in Search Markets: An Empirical Case Study

Tobias Salz

Massachusetts Institute of Technology

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 who do not use intermediaries. I investigate the magnitude of these effects in New York City's trade-waste market, where buyers can search either 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.

I. 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,

I am indebted to my advisors John Asker, Guillaume Frech  tte, 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 Kalouptsi, Yunmi Kong, John Lazarev, Konrad Menzel, Joseph Mullins, Ariel Pakes, Maryam Saeedi, Maher Said, Kevin Thom, Allan Collard-Wexler, Larry White, Sevgi Yuksel, and Jidong Zhou as well as seminar participants at Caltech, Carnegie Mellon University, Columbia, Cornell, Duke, Harvard, London School of Economics, Microsoft Research Cambridge, Massachusetts Institute of Technology, Northwestern, the Federal Reserve Bank of New York, New York University, Princeton, Stanford, University of California Los Angeles, University of Southern California, University of Pennsylvania, University of Texas at Austin, and Yale School of Management provided invaluable feedback. The research in this paper was conducted with data provided by New York City's Business Integrity

Electronically published December 14, 2021

Journal of Political Economy, volume 130, number 2, February 2022.

   2021 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/717349>

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 data set from the New York City trade waste industry to 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 who 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 because of the very nature of these markets. However, the trade waste industry in New York provides 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.

Commission. The data have 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. This paper was edited by Ali Hortaçsu. Replication files are provided as supplementary material online.

¹ Spulber (1996b) estimates that 25% of US economic activity is intermediation. Examples beyond housing and real estate include textile brokers, advertising brokers, freight brokers, literary agents, head hunters, and energy brokers.

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

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 delegates 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 (e.g., Hortagsu and Syverson 2004; Hong and Shum 2006). 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 nonbrokered 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. However, 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%.

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 who 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 who 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 who 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 is a 14.2% increase.³

A. 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 who 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 (e.g., Rubinstein and Wolinsky 1987; Gehrig 1993; Spulber 1996b; Lizzeri 1999; Rust and Hall 2003; Moraga-González, Sándor, and Wildenbeest 2014). A closely related study is Gavazza (2016), who examines the effect of intermediaries in the secondary market for business aircraft using a dynamic search and bargaining framework. Another paper that compares a bilateral market with an intermediated market is Hendel, Nevo, and Ortalo-Magné (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 nonsequential search setting, respectively. The corresponding equilibrium models for sequential and nonsequential search settings were formulated in Burdett and Judd (1983) and Stahl (1989). Goods are homogeneous in these models and price dispersion arises because of mixed strategy pricing. For the purposes of this study, the assumption of nonsequential 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 facilitates identification. Third, several studies have found that nonsequential search better explains actual search behavior (De Los Santos, Hortaçsu, and Wildenbeest 2012; Honka and Chintagunta 2014).

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.

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 study that 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 for both the sequential and the nonsequential cases. Allen, Clark, and Houde (2014) use Canadian mortgage data along with quasi-experimental variation due to a merger to estimate search cost nonparametrically. With the exception of Allen, Clark, and Houde (2014, 2019), 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.

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. The pricing subgame of sellers in the search market can be viewed as a first-price auction with an unknown number of competitors, which depends on the customer search strategy. The identification of auction models has been discussed in Guerre, Perrigne, and Vuong (2000) and further developed in Athey and Haile (2002) for asymmetric auctions.⁵

Roadmap.—Section II provides relevant industry facts and describes the data. Section III establishes important descriptive facts that inform the setup of the model. Section IV describes the model, section V describes the identification of the model, and section VI describes the estimation. Section VII describes the results, and section VIII describes the counterfactual computations.

⁴ Factors that determine the cost of service provision include the location, quantity, and composition of the waste; 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), time incentives in procurement projects (Lewis and Bajari 2011), collusion in auctions (Asker 2010), bid preference programs (Krasnokutskaya and Seim 2011), and many others.

II. 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 V.

A. Data

Trade waste industry is the official name for New York's private waste market.

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 (BIC), has a comprehensive oversight mandate over the trade waste industry.⁶

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 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.

Table 1 provides 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 pickups per week is five. Close to half of all businesses generate recyclables.⁷ One data caveat is that I observe only 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

⁶ 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 June 2013 (which is the relevant data period), and \$18.27 from July 2013 onward. If measured per 100 lb, it was \$8.00 before 2008, \$10.41 from 2008 to June 2013, and \$11.98 thereafter. Another important restriction regards the length of contracts, which cannot exceed 2 years, after which the customer has the option to sign with another carter.

⁷ Contracts differ in whether the price is charged per cubic yards or per pounds. For my analysis, I use only 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. I also exclude contracts that involve only medical waste, shredding of paper, and cardboard or grease haul.

TABLE 1
SUMMARY STATISTICS

Variable	Mean	Median	Standard Deviation
Monthly charges (\$)	198	92	239.2
Price (\$)	11.93	12.2	3.19
Monthly quantity (cubic yards)	25.94	8.66	55.5
Recyclables (yes/no)	.5	.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.

set of observables that are available in the broker data set as well as in the customer register and impute the broker fee for a given contract as the predicted value from this regression (see sec. VI for details).

B. 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 carter. A breakdown by business type is provided in 1 in appendix E. Carter service consists of the pickup and hauling of recyclable and nonrecyclable refuse to one of the 61 transfer stations in or around New York (indicated by triangles in fig. 1).⁸

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. 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 10 biggest firms 55% of all customers.⁹

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 (fig. 1), which is more than one-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.¹⁰

⁸ The typical volume of one of the rear loaders used in New York City is about 20–30 cubic yards, and the median monthly quantity generated by a business in the city is about 8.66 cubic yards. 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.

⁹ For confidentiality reasons, I cannot provide a more complete picture of the firm-size distribution.

¹⁰ 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, e.g., Jacobs, Friel, and Raddick 2001).

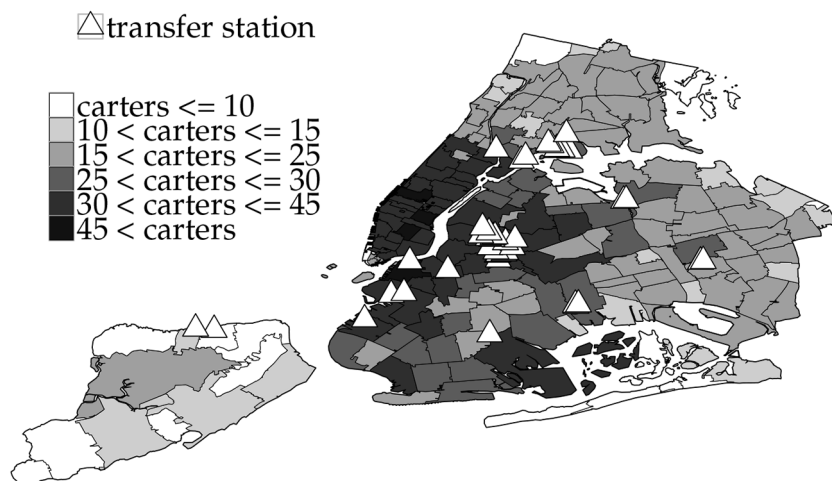


FIG. 1.—Average number of active carters per zip code averaged over time. Active is defined as having at least one customer in the zip code. Triangles show the locations of the transfer stations.

The following observation from an article in the *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. (Keenan 2014)

When the waste industry was consolidating on a national scale, New York City was still in 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 overbillings of 30%–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 why 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 with other cities or rural areas.

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 (see, e.g., Rosengren 2015). 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.

C. Brokers

Trade waste brokers are a potential remedy for customers' search problem, and their services are particularly valuable in markets like New York's. 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.¹¹ Conversations with brokers reveal that they award contracts through a *request for proposals*—a competitive bidding process akin to a first price-auction—to these carters.¹²

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 volumes 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

¹¹ Brokers arrange contracts between businesses and carters; they are not allowed to accept direct payments by carters (see <https://codelibrary.amlegal.com/codes/newyorkcity/latest/NYCAAdmin/0-0-0-27182>).

¹² 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. A 2014 article in the *New York Times*, which portrays one of the large trade waste brokers, reiterates some of these points: "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" (Zax 2014).

the three remaining quartiles (15%, 14%, and 19% brokered contracts).¹³ Appendix table 2 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 hotels and institutions are more likely to use brokers than other nonfood retail and wholesale businesses.¹⁴

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 10 firms for 52.6%, and the largest 15 firms for about 78.95% averaged over the entire observation period. Note, however, that brokers operate across cities and that market shares in New York City might not be the most meaningful.

III. 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.

A. Evidence of Price Dispersion

I document both variation in raw prices p_{ijt} and, following Allen, Clark, and Houde (2014), the dispersion in residual prices to account for observable price variation. Residual prices \hat{p}_{ijt} are computed from the following regression: $p_{ijt} = \mathbf{X}_{ijt} \cdot \boldsymbol{\beta} + \hat{p}_{ijt}$. I restrict the sample to customers who are in the retail business to limit the possibility that observed price variation is driven by unobserved differences across customer 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 cubic yard) 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 nonbrokered contracts and include a common set of controls; the only difference is

¹³ These numbers pertain to the subset of contracts that are used in the study. In app. A, I explain how I select the sample.

¹⁴ My model assumes that substitution between brokers and the search market is solely based on search cost. This modeling decision is supported by the observation that almost all carters who are active in the broker market also serve contracts in the regular market.

TABLE 2
DOCUMENTING PRICE DISPERSION

	FIRST SPECIFICATION (No Carter Fixed Effects)		SECOND SPECIFICATION (Carter Fixed Effects)	
	Not Brokered	Brokered	Not Brokered	Brokered
$1 - R^2$.71	.35	.51	.24
SD (p_{ijt})	2.9	4.4	2.9	4.4
SD (\tilde{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 \tilde{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, and fixed effects for the number of pickups per week.

whether I include carter fixed effects. Throughout, I use q to denote the quantity of waste produced by a given customer.

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 is 0.24 in the search market and 0.2 in the broker market. 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.¹⁵

¹⁵ The following calculations give a sense of the dollar value of this dispersion: for a business that reports in cubic yards and generates the mean quantity of waste (17.8 cubic yards), moving 1 standard deviation (SD (\tilde{p}_{ijt})) in the residual distribution of prices would cost an extra \$1,281.6 over the length of a contract, which is slightly longer than 2 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 10 employees. BizBuySell provides the median sales price of a business in New York City that has been sold over this platform, which was about \$229,500 in 2013. If we assume an interest rate of 3%, that would imply annual profits of \$6,885.0. The subset of businesses sold on this page is almost surely biased toward small companies, but it provides some reference point for the above calculations.

*B. Comparing Prices across Broker Market
and Search Market*

This section establishes that broker prices p^B including the broker commissions ϕ are higher than search market prices p^S , which are in turn higher than prices in the broker market without commissions:

$$p^B \cdot (1 + \phi) > p^S > p^B.$$

One explanation for the surprising fact that buyers who use brokers end up paying higher prices is that they have higher search costs. I establish this result through two sets of regressions (shown in appendix tables 3 and 4), both of which have prices as the dependent variable. In the first set, brokered contracts include only the broker price p^B , whereas in the second set, brokered contracts also include the commission $p^B \cdot (1 + \phi)$. The main variable of interest is a dummy that indicates whether the contract is brokered. I explore both the mean effect in an ordinary least squares 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 zip code Herfindahl-Hirschman index. The estimated linear effect of the broker dummy from the first specification in appendix table 3 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.

Appendix table 4 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.

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.

IV. Model Overview

To evaluate the effect of brokers, one needs to take into consideration the selection of customers into the broker market and the resulting pricing incentives of sellers in both markets. The model will address these challenges.

The search of businesses for brokers is modeled as a sequential game between customers and carters. Brokers are nonstrategic players. A customer j in the model is described by privately observed independently and identically distributed (i.i.d.) search cost κ_j , drawn 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 i.i.d. 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. At $t = 0$, the customer draws search cost κ and carters draw service cost c . Both are privately observed. At $t = 1$, customers decide whether to delegate search to a broker and, if not, how many price quotes to get, $m \in \{1, \dots, M\}$, where M is the total number of carters. At $t = 2$, carters submit price quotes either in a first-price auction when the contract is procured through a broker, or in the search market.

A. 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 on the basis of 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 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 section IV.B, I will provide details on the bidding functions that give rise to $\mathbb{E}[p^B]$ and $\mathbb{E}[p^{1:m(\kappa)}]$. The markup charged by brokers is denoted as ϕ .

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

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

Every customer with a 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 $c_{ij} \sim \mathcal{G}(\cdot)$ is the distribution of cost draws for carters.¹⁶ I will now discuss what the optimal search strategy of customers against such an equilibrium price offer function looks like. Search is nonsequential.

A customer j with search cost κ_j minimizes their total cost over the number of searches $m \in \{1, \dots, M\}$:

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

if one uses 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

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

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.

LEMMA 1. There are marginal types $0 \leq \kappa_{M-1} < \dots < \kappa_{m-1} < \kappa_m < \bar{\kappa} \leq \infty$ 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.

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 cutoff

¹⁶ The distribution $\mathcal{F}(\cdot)$ is unobserved. The data record only the contract price.




FIG. 2.—Sorting of buyers according to their search cost. Higher search cost leads to fewer calls m . Types with search costs above $\bar{\kappa}$ delegate their search to a broker.

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 $q \cdot \mathbb{E}[p^B] \cdot \phi > \mathbb{E}[p^{1:m}] - \kappa \cdot m$, then $q \cdot \mathbb{E}[p^B] \cdot \phi > q \cdot \mathbb{E}[p^{1:(m+1)}] - \kappa \cdot (m+1)$. QED

Figure 2 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.

B. 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 the broker market are linked through the common cost distribution $\mathcal{G}(\cdot)$ that carters in both markets are drawing from. Denote $\tilde{\mathcal{G}}(\cdot) = 1 - \mathcal{G}(\cdot)$. 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:

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

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 cutoff types. I will now describe carter's strategy $\beta_s(\cdot)$ in the subgame 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. With this notation, the maximization problem for carters is

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

This maximization problem is akin to a first-price procurement auction with an unknown number of competitors. The number of competing firms is determined by customers' search strategies, as summarized by the weights. Suppressing the dependence on covariates, it can be shown that the bidding function from this problem has a closed form solution (Krishna 2009) and is given by

$$\beta_s(c) = \sum_{m=1}^M \left[\frac{w_m \cdot \tilde{\mathcal{G}}(c)^{(m-1)}}{\sum_{k=1}^M w_k \cdot \tilde{\mathcal{G}}(c)^{(k-1)}} \cdot \left(c + \frac{1}{\tilde{\mathcal{G}}(c)^{(m-1)}} \int_c^{\tilde{c}} \tilde{\mathcal{G}}(u)^{(m-1)} du \right) \right]. \quad (3)$$

Combining the behavior of carters and the customer search strategy, we find that 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: $\beta_b(\cdot)$, $b \in \{1, \dots, B\}$;
2. a bidding strategy in the search market: $\beta_s(\cdot)$; and
3. customer search weights w_1, \dots, 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. $\beta_b(\cdot)$, $b \in \{1, \dots, B\}$ 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, \dots, w_M ; and
3. $\mathcal{F}(\cdot)$ and $\mathcal{F}^B(\cdot)$ result from $\beta_s(\cdot)$ and $\beta_b(\cdot)$, $b \in \{1, \dots, B\}$.

V. 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 am only able to use a proxy for the number of bidders, which is the number of carters that a broker awards contracts to. This issue is discussed in more detail in section VI.B.2, and I also conduct a robustness check that explores the sensitivity of the counterfactual results to changes in the number of bidders.

From Athey and Haile (2002), I adopt the following definition of identification. Define a model as a pair (\mathbb{F}, Γ) , where \mathbb{F} is a set of joint distributions over the vector of latent random variables, Γ is a collection of mappings $\gamma: \mathbb{F} \rightarrow \mathbb{H}$, 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 (\mathcal{T}, γ) generating the observables.

DEFINITION 2. A model (\mathbb{F}, Γ) is identified iff for every $(F, \hat{F}) \in \mathbb{F}^2$ and $(\gamma, \hat{\gamma}) \in \Gamma^2$, $\gamma(F) = \hat{\gamma}(\hat{F})$ implies $(F, \gamma) = (\hat{F}, \hat{\gamma})$.¹⁷

A graphical illustration of simulation results 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 on both search costs via the weights and carter costs. The simple intuition for nonidentification is that the same distribution of prices can be generated either by a combination of low search cost and high carter cost or by a combination of high search cost and low carter cost:

$$\mathcal{F}^o(p) = \sum_m w_m \cdot (1 - [1 - \mathcal{G}(\beta_s^{-1}(p))]^m). \quad (4)$$

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 $g(\cdot)$ and where $(1 - \mathcal{G}(\cdot))/g(\cdot)$ is Lipschitz continuous.¹⁸

Furthermore, I require the following assumption.

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

¹⁷ For a similar definition, see Matzkin (2007).

¹⁸ In an early version of this paper, the model accounted for persistent differences of carters by allowing for a finite number of different cost distributions. In app. B.3, I discuss the assumptions under the previous setup and briefly outline how one can estimate such a model.

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, Perrigne, and Vuong (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 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 fig. 2) 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 assumptions 1 and 2 hold, the search weights (w_1, \dots, w_M) and cutoff types $(\kappa_1, \dots, \kappa_{M-1}, \bar{\kappa})$ 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. (1) 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 $\beta_s(\cdot)$. (2) Given that $\mathcal{G}(\cdot)$ and $\beta_s(\cdot)$ are known, the weights are uniquely identified on the basis of an inductive argument. (3) 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 cutoff types $\kappa_1, \dots, \kappa_{M-1}$ are a direct function of those conditional prices and q and therefore also identified. Last, 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 2, 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 $\kappa_{M-1}, \dots, \kappa_1, \bar{\kappa}$. 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. Appendix figure 5a provides an example based on my estimates. Each line in those graphs has a corresponding known point on the empirical cumulative distribution function (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 cutoff types $\kappa_m = q \cdot (\mathbb{E}[p^{1:m}] - \mathbb{E}[p^{1:(m+1)}])$, the identified partition becomes much finer. Appendix

figure 5*b* shows this partition for the empirical implementation of the model, which is based on four different quantity levels.¹⁹

VI. Estimation

This section explains the details of the estimation. Estimation is conducted via a 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. Last, 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 \mathbf{x} . The contract covariates \mathbf{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. Conditioning on quantity accounts for potential scale efficiencies. For search cost, I estimate an unconditional distribution.²⁰

Regarding the length of the contract, I rely on the fact that the BIC mandates a maximal contract lock-in of 2 years.²¹ The estimation assumes that when customers search, they expect to contract for the full 2 years. I model only the initial 2 years and do not try to model the dynamics of contract renewals. Hence, in the estimation I rely on only the initial price in the customer-carter relationship. One issue is that some customers might anticipate that they remain in contract for only a few months and therefore have lower incentives to search. To better fit my assumption of a 2-year planning horizon, I use contracts only for customers who stay in the contract for 2 years. The data show that 58.4% of all contracts last the full 2 years, 10% for 1.5 years, 13% for 1 year, and 17.8% for only 6 months. I provide an in-depth discussion of dynamic considerations in the customer-carter relationship in section VI.B.1.

I further restrict the estimation to a subset of contracts whose customers operate a retail business to minimize the possibility that price

¹⁹ 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 cutoff types $\kappa_1, \dots, \kappa_{M-1}$ to the right, but also it 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 sec. VI, where I will also allow the substitution to brokers to depend on q .

²⁰ 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.

²¹ See <https://www1.nyc.gov/site/bic/industries/customer-information.page>. 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.

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 therefore does not restrict the sample too severely.²²

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 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

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

The specification for the customer search cost distribution is²³

$$\mathcal{H}(\cdot) = \mathcal{N}_{[0,\infty)}(\mu^s, \sigma^s). \quad (6)$$

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 because of the restrictions that changes in the quantity impose on buyers' search strategy and their type distribution.²⁴ Hence, I assume that the relative attractiveness of brokers also depends linearly on q via ψ as follows:

$$q \cdot \mathbb{E}[p^B|x] \cdot \phi(x) - \psi \cdot q = q \cdot \mathbb{E}[p^{1:m(\bar{x})}|x] + m(\bar{\kappa}(x)) \cdot \bar{\kappa}(x). \quad (7)$$

²² See appendix table 1 for a breakdown of customer business types. These restrictions leave 44,417 contracts on which the model is estimated. Some additional details on the construction of the sample are given in app. A.

²³ I have also tried a version where customer search cost depends on the borough, but this made almost no difference in the results, which is why I decided to drop the extra parameters. The difference in average prices across boroughs is less than \$0.50.

²⁴ I illustrate this issue in app. D.1. The model predicts that higher-quantity buyers search more, which leads all cutoff types, including the broker marginal type, to shift rightward. Holding the search 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. Appendix table 2 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 . An alternative explanation is that high-quantity buyers are compensated in other ways, such as additional services from brokers, or that large-quantity buyers command lower commissions in ways that is not directly captured by my data. If the latter is true, one could interpret ψ as a reduced form way of capturing the increased bargaining power for lower commissions that higher-quantity buyers have. The first explanation has less appeal because search costs are interpreted as a marginal cost.

A final data caveat is that I have access to broker commissions for only 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 data sets, I run a hedonic pricing regression to impute the broker commissions ϕ based on the two sets of variables that overlap in both data sets, 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 reveal 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 $\mathbb{E}[p^B|x] \cdot (1 + \bar{\phi}(x))$. 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 $\theta = \{\mu^s, \mu^c, \sigma^s, \sigma^c, \gamma^{s,Boro}, \gamma^r, \gamma^q, \psi\}$. 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 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 distribution of quantity q , which is the only continuous variable in x , into its four quartiles. The product of four 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 candidate parameter. 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 $\beta_b(\cdot) \forall b, \beta_s(\cdot), w_1, \dots, w_M, \kappa_1, \dots, \kappa_{M-1}, \bar{\kappa}$. I match five different types of moments for each of the aforementioned cells x . These moments are based on the mean and the standard deviation of prices for both brokered and nonbrokered contracts as well as the fraction of brokered contracts. Note that the model-derived mean of prices in the search market is given by

$$\Upsilon_s(\theta, x) = \sum_{m=1}^M w_m(x) \cdot \int \beta_s(c|x) \cdot g_{1:m}(c|x; \theta) dc,$$

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 $\{w_1, \dots, w_M\}$ 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 $\{p_{s_1}, \dots, p_{s_K}\}|x$ 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 $p_{s_b}^b = \beta_b(\underline{c}_{s_b}^k; N_b)$:

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

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. Last, 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

$$m_{1,s}(\theta, x) = N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} p_i - K^{-1} \cdot \sum_{k=1}^K p_{s_k}(\theta, x), \quad (8)$$

and the moments for the standard deviation of prices are given by

$$\begin{aligned} m_{2,s}(\theta, x) = & \left(N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} (p_i - N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} p_i)^2 \right)^{0.5} \\ & - \left(K^{-1} \cdot \sum_{k=1}^K (p_{s_k}(\theta, x) - K^{-1} \cdot \sum_{k=1}^K p_{s_k}(\theta, x))^2 \right)^{0.5}. \end{aligned} \quad (9)$$

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 $\{p_{s_1}^b, \dots, p_{s_K}^b\}|x$. Last, the moment for the fraction of brokered contracts is given by

$$m_f(\theta, x) = N_x^{-1} \cdot \sum_{i \in \mathcal{A}_x} \mathbb{1}\{\text{brokered}\}_i - (1 - \mathcal{H}(\bar{\kappa}(x))) \cdot N_x. \quad (10)$$

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

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

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

$$\hat{\theta}_{MSM} = \operatorname{argmin}_{\theta} \mathbf{m}(\theta)' \cdot \Omega \cdot \mathbf{m}(\theta). \quad (11)$$

I use the inverse of the observed variance of each moment in the data to weight the moments.²⁵ I employ a nonparametric bootstrap (Efron and Tibshirani 1994) to obtain standard errors and 95% confidence intervals for my estimates and the counterfactual numbers. I draw 400 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.

A. The Estimation Algorithm

For each parameter guess, constructing the moments requires solving for the equilibrium objects $\beta_b(\cdot) \forall b, \beta_s(\cdot), w_1, \dots, w_M, \kappa_1, \dots, \kappa_{M-1}, \bar{\kappa}$. 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 cutoff 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 who 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 (Diamond 1971)—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.²⁶ This drastically speeds up the inner loop because I do not have to run a separate numerical search for the equilibrium.

One potential concern is that the set of equilibria with a nondegenerate price distribution is not a singleton. 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

²⁵ For each moment, I use 20 simulation draws, which leads to a total of 1,280 simulation draws.

²⁶ The Diamond paradox is rejected by the data.

TABLE 3
ESTIMATION ALGORITHM

Algorithm 1: Estimation of Model

Result: Estimate of θ

```

while  $\mathbf{m}(\theta)' \cdot \hat{\Omega} \cdot \mathbf{m}(\theta) > \text{outer tolerance}$  do
  1. Initialize weight vector  $w_0$  with  $1/M$  in each entry;
  2. while  $d(w^k, w^{k-1}) = \|w^k - w^{k-1}\| > \text{inner tolerance}$  do
    2.1 Recompute the broker bidding functions  $\beta_b^k(\cdot|\theta) \forall b$ ;
    2.2 Use  $w^k$  to recompute the bidding function  $\beta_s^k(\cdot|w^k; \theta)$ ;
    2.3 Use  $\beta_s^k(\cdot|w^k; \theta)$  to compute expected prices  $E[p^{1:m}]$  for each  $m$ ;
    2.4 Recompute  $\kappa_m = E[p^{1:m}] - E[p^{1:m+1}] \forall m \in 1, \dots, M$  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  $\{p_{\kappa}, \dots, p_{\kappa}\}$  and  $\{p_{\kappa}^b, \dots, p_{\kappa}^b\}$ , compute the
    average and the standard deviation of simulated prices, compute fraction of bro-
    kered contracts  $(1 - \mathcal{H}(\bar{\kappa})) \cdot N$ ;
  4. Construct moments for the objective function.
end

```

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. Appendix G shows the results from these explorations. For a wide variety of starting values, the algorithm converges to the same equilibrium in a few steps.

B. Discussion of Model Limitations

1. Contract Dynamics

The estimation focuses on new contracts between buyers and carters. The welfare quantifications that I discuss below pertain to only 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 provide 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 regard is the BIC regulation that customers can switch after 2 years at no additional pecuniary cost. I therefore assume that each first-time relationship lasts for 2 years. In addition, to rule out unobserved heterogeneity, I include only contracts in the estimation that lasted for two full years and where the customer did not either exit or switch carters.²⁷

²⁷ What makes this less problematic is that similar to q , contract length scales the cutoff 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 cutoff types κ and therefore the estimated search

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 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.²⁸

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 who 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; 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.

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

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.

²⁸ 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. To give a sense of how large the dynamic component of pricing is, fig. 3 in app. 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 2 years, prices adjust upward, consistent with search costs acting as switching costs. However, the magnitudes, at less than \$0.50, 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 is left for future research.

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 (see, e.g., Asker 2010; Krasnokutskaya 2011). 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, Hong, and Shum 2003).

While it is outside the scope of this paper to extend 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 misallocation. This means that any allocative effects due to cost differences in the counterfactual are overstated if some 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 costs are underestimated. Because brokers increase allocative cost efficiencies less when costs are less dispersed but reduce search cost more when search costs are higher, the overall bias on the counterfactual welfare results is unclear.

VII. Results

This section describes the estimation results. I show graphs of 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 costs per cubic yard are estimated as \$9.97, 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 negative effect of recyclables on cost, but the coefficient is not significant.

TABLE 4
MODEL PARAMETER ESTIMATES

Parameter	Estimate	SE	95% CI
Supply:			
Carter cost (\$):			
Mean	9.969	.24	9.776–10.023
SD	2.96	.157	2.81–3.105
Cost efficiencies ($\times 1,000$):			
Quantity	–2.456	1.709	–4.62 to .259
Recyclables	–.152	7.048	–1.589 to 19.8
Cost shifter:			
Bronx	–.235	.031	–.3 to –.214
Brooklyn	–.002	.007	–.003 to .022
Manhattan	–.32	.033	–.398 to –.276
Demand:			
Search cost (\$):			
Mean	79.718	5.298	77.302–95.007
SD	62.352	4.903	58.929–73.238
Quantity broker shift	–.309	.015	–.333 to –.284

NOTE.—The table shows the parameter estimates along with bootstrapped standard errors and 95% confidence intervals (CIs), based on 400 bootstrap iterations. The borough cost shifters are relative to Queens.

There are small cost differences across boroughs. The Bronx and Manhattan are estimated to be cheaper to service than Brooklyn and Queens.

The mean search cost estimate is \$79.7, with a large standard deviation of \$62.4. The coefficient that captures differential selection to brokers due to differences in quantity is estimated as -0.31 . This means that with each additional cubic yard, 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 pecuniary incentives to do so.

Table 5 shows the average cost (per cubic yard) 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 costs. This is mostly because buyers with higher q have higher incentives to search, leading the cost to be more selected. Figure 3 shows this effect graphically. In table 5, $Q_{a,b}$ refers to quantities from quantile a to b .

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, 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 who use brokers and those who do not gives a sense of the selection into broker services. The search costs for customers who use brokers are about two to three times as high as those who contact carters directly. The estimates

TABLE 5
EXPECTED COST (\$) FOR CARTER PER CUBIC YARD PER MONTH

	Bronx	Brooklyn	Manhattan	Queens
$Q_{0,25}$	9.8 (.687)	11.9 (.71)	9.5 (.657)	9.3 (.724)
$Q_{25,50}$	8.9 (.733)	9.2 (.764)	9.5 (.729)	9.2 (.736)
$Q_{50,75}$	7.9 (.656)	6.9 (.731)	7.8 (.657)	7.9 (.705)
$Q_{75,100}$	5.4 (.496)	5.7 (.433)	5.5 (.442)	5.7 (.466)

NOTE.—The table shows the average cost for different boroughs. To compute this cost, I draw repeatedly (10,000 times) from the distribution of bidders. $Q_{a,b}$ refers to quantities from quantiles a to b . Bootstrapped standard errors are in parentheses (400 iterations).

therefore reflect the high markups charged by brokers. The wedge that this markup generates between the broker market and the nonbrokered market is rationalized through higher search costs for buyers who use broker services.

Table 6 also gives an overview of the total costs that customers incur for searching—that is, 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




FIG. 3.—Population distribution of carter cost and realized distribution for two different equilibrium search weights (2.5 and 9.5 inquiries on average).

TABLE 6
SEARCH COST AND INQUIRIES

SUBSET	SEARCH COST PER INQUIRY (\$)		NUMBER OF SEARCHES $\kappa < \bar{\kappa}$ (Searches)	TOTAL SEARCH COST (\$) $\kappa < \bar{\kappa}$ (Searches)	FRACTION OF TOTAL EXPENSES
	$\kappa < \bar{\kappa}$ (Searches)	$\kappa > \bar{\kappa}$ (Brokers)			
Bronx:					
$Q_{0,25}$	87.3 (10.092)	220.1 (25.508)	1.2 (.14)	93.0 (10.738)	.15
$Q_{25,50}$	81.5 (9.478)	194.7 (22.746)	1.6 (.184)	99.1 (11.446)	.08
$Q_{50,75}$	83.2 (9.745)	200.7 (23.912)	2.9 (.337)	163.0 (18.813)	.08
$Q_{75,100}$	72.1 (8.634)	170.3 (20.263)	8.2 (.956)	536.5 (62.219)	.09
Brooklyn:					
$Q_{0,25}$	85.6 (9.877)	210.7 (24.302)	1.2 (.14)	92.9 (10.727)	.15
$Q_{25,50}$	75.2 (8.666)	177.1 (20.416)	1.7 (.191)	98.4 (11.356)	.08
$Q_{50,75}$	57.4 (7.269)	145.3 (17.583)	3.8 (.451)	165.7 (19.13)	.08
$Q_{75,100}$	63.3 (7.522)	154.2 (17.99)	8.7 (1.011)	510.9 (59.18)	.08
Manhattan:					
$Q_{0,25}$	87.8 (10.152)	223.6 (25.908)	1.2 (.14)	93.0 (10.741)	.15
$Q_{25,50}$	83.4 (9.683)	201.5 (23.528)	1.6 (.182)	99.3 (11.466)	.08
$Q_{50,75}$	85.8 (9.962)	212.0 (24.926)	2.8 (.329)	161.9 (18.688)	.08
$Q_{75,100}$	75.8 (9.0)	178.6 (21.215)	8.0 (.935)	545.8 (63.143)	.09
Queens:					
$Q_{0,25}$	85.6 (9.88)	210.6 (24.316)	1.2 (.14)	92.9 (10.727)	.15
$Q_{25,50}$	75.1 (8.675)	177.0 (20.437)	1.7 (.191)	98.4 (11.357)	.08
$Q_{50,75}$	57.4 (7.073)	145.3 (17.331)	3.8 (.449)	165.7 (19.131)	.08
$Q_{75,100}$	63.2 (7.533)	154.0 (18.016)	8.7 (1.009)	510.6 (59.242)	.08

NOTE.—The table shows expected search cost per inquiry, number of inquiries, and total expenses for search. Bootstrapped standard errors are in parentheses.

therefore removes the ability to search but also the cost of searching. A full welfare evaluation should take the cost of searching into account.²⁹

VIII. 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. More buyers are now searching, which changes the composition of search costs. Buyer’s overall expenses are rising. Brokers lose commissions, and carters’ expected profits change

²⁹ I have also made an effort to compare my estimates with the literature. Allen, Clark, and Houde (2019) compare the search cost from different models (Hortaçsu and Syverson 2004; Hong and Shum 2006; Honka 2014). 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. From 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 nonsequential 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 ranges between 10% for online transactions and 35% for offline transactions. In my setting, I find a median ration of 26% and a mean of 34%. This is broadly in line with the aforementioned papers.




FIG. 4.—Cutoff types and counterfactual cutoff types. This figure shows an example of a search cost density along with the old and new marginal κ types.

because of the changing composition of buyers. Welfare declines because of larger expenses for search and a higher average contract cost.

Figures 4 and 5 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 4 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

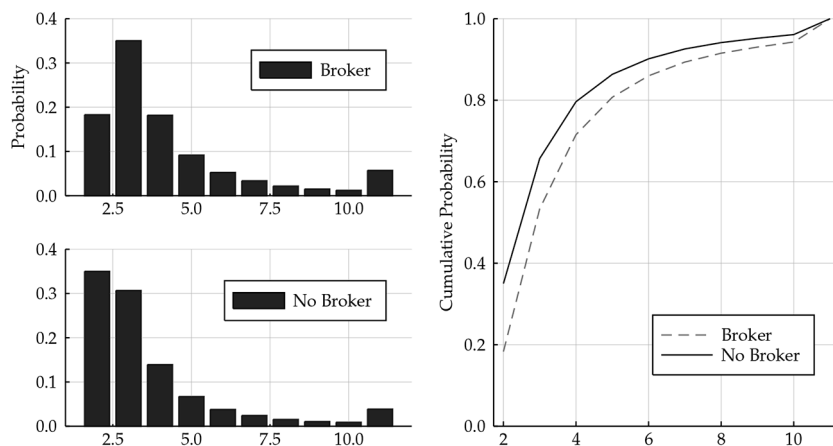


FIG. 5.—CDF of equilibrium search strategies. This figure shows the cumulative CDFs for the original and the counterfactual search strategy (*right*) and the two histograms (*left*).

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 cutoff type (“New Marginal Types”). 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 who were already in the search market to also search less. This effect compounds the selection effect. In figure 4, 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 compares the histogram of price inquiries with and without brokers and shows that everybody to the left of the “Broker Marginal” falls into the bin of two price inquiries. 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. Appendix table 5 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 in the baseline that are unrelated to search cost, as measured by ψ . These buyers now have strong incentives to search in the counterfactual, leading to more price draws than

TABLE 7
COUNTERFACTUAL OVERVIEW

	CHANGE IN BUYER EXPENSES			CARTER		WELFARE (Total Cost)	
	Not Brokered	Brokered	All	Margin	Profits	Lower Bound	Upper Bound
Δ absolute (\$)	64.0	445.0	127.0	.046	−11.1	4.28	12.61
SE	(11.0)	(157.0)	(32.0)	(.0271)	(10.0)	(1.1)	(3.07)
95% CI	47.7–80.2	265.0–759.3	88.3–189.0	.002–.088	−28.7 to 1.4	3.28–6.78	9.63–18.77
Δ percent	2.52	11.7	4.6	1.95	−1.8	4.41	14.22
SE	(.57)	(1.96)	(1.19)	(1.14)	(1.53)	(1.2)	(4.03)
95% CI	1.84–3.52	8.9–15.42	3.17–6.95	.1–3.81	−4.53 to .25	3.35–7.19	10.58–22.77

NOTE.—The table shows expected search cost per inquiry, number of inquiries, and 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 (in parentheses) and confidence intervals (CIs) are based on 400 iterations.

brokers are soliciting through the auctions. Appendix E shows that the carter profit decrease is coming mainly from buyers with very high quantity, who are responsible for an disproportionate percentage of profits. Carter profits and the welfare effects therefore depend on both the search cost and the quantity of buyers, which together determine how much counterfactual search occurs.

Buyer expenses include both the contract and the search expenses. Eliminating brokers for buyers who 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 who benefit from the externality in the current market setting is much larger than the number of buyers who 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 at \$2.4 nearly as large. 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 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 is \$4.3 million. This would imply a 4.4% decrease in welfare at the lower end and a 14.2% decrease at the higher bound.

The welfare loss is a combination of higher search cost and a reallocation effect. The reallocation effect means that without brokers, the market engages in less search and picks carters who 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 cubic yard is 2.3% higher without brokers.

A. Robustness to Number of Bidders in the Auction

One limitation of my analysis is that the number of bidders on a broker contract is inferred from the number of carter relationships that each broker holds. 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 who bid but never win. To probe the sensitivity of the counterfactual results, I reestimate 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 appendix table 7.

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 nonbrokered 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 one-third, 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 that while there are changes in quantitative magnitudes, those changes are modest.

B. 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. In general, these services are therefore underprovided since intermediaries create social returns that are not reflected in their private payoffs.³⁰ While this result provides a rationale for policies that promote intermediation, it also highlights an important insight for regulatory information provision and disclosure policies (e.g., see Jin and Leslie 2003): it means that such policies can have large effects due to supply-side responses, even if they reach only 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, such as for investment goods, is that sellers are very dispersed around the globe and products are ordered with idiosyncratic specifications. One would therefore expect that search frictions are severe

³⁰ The efficiency of broker subsidies will also depend on potential duplication of fixed cost if they affect the extensive margin (Mankiw and Whinston 1986).

in these markets and, therefore, that the framework in this paper can be applied to them.

IX. 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 it more competitive for sellers. To quantify these effects of intermediation, I use a new and detailed data set 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 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. 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 inquiries. The latter is then used to recover buyers' search costs. The estimates reveal that buyers' search costs are an important factor to take into account for policy considerations, such as the comparison between the current 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. 2014. "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry." *A.E.R.* 104 (10): 3365–96.
- . 2019. "Search Frictions and Market Power in Negotiated-Price Markets." *J.P.E.* 127 (4): 1550–98.
- Armstrong, Mark. 2015. "Search and Ripoff Externalities." *Rev. Indus. Org.* 47 (3): 273–302.
- Asker, John. 2010. "A Study of the Internal Organization of a Bidding Cartel." *A.E.R.* 100 (3): 724–62.
- Athey, Susan, and Philip A. Haile. 2002. "Identification of Standard Auction Models." *Econometrica* 70 (6): 2107–40.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–69.
- De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *A.E.R.* 102 (6): 2955–80.
- Diamond, Peter A. 1971. "A Model of Price Adjustment." *J. Econ. Theory* 3 (2): 156–68.
- Efron, Bradley, and Robert J. Tibshirani. 1994. *An Introduction to the Bootstrap*. Boca Raton, FL: CRC.
- Gavazza, Alessandro. 2016. "An Empirical Equilibrium Model of a Decentralized Asset Market." *Econometrica* 84 (5): 1755–98.
- Gehrig, Thomas. 1993. "Intermediation in Search Markets." *J. Econ. and Management Strategy* 2 (1): 97–120.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica* 68 (3): 525–74.
- Haile, Philip A. 2001. "Auctions with Resale Markets: An Application to US Forest Service Timber Sales." *A.E.R.* 91 (3): 399–427.
- Haile, Philip A., Han Hong, and Matthew Shum. 2003. "Nonparametric Tests for Common Values at First-Price Sealed-Bid Auctions." Working Paper no. 10105, NBER, Cambridge, MA.
- Hendel, Igal, Aviv Nevo, and François Ortalo-Magné. 2009. "The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com." *A.E.R.* 99 (5): 1878–98.
- Hong, Han, and Matthew Shum. 2006. "Using Price Distributions to Estimate Search Costs." *RAND J. Econ.* 37 (2): 257–75.
- Honka, Elisabeth. 2014. "Quantifying Search and Switching Costs in the US Auto Insurance Industry." *RAND J. Econ.* 45 (4): 847–84.
- Honka, Elisabeth, and Pradeep K. Chintagunta. 2014. "Simultaneous or Sequential? Search Strategies in the US Auto Insurance Industry." Working paper.
- Hortaçsu, Ali, and Chad Syverson. 2004. "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds." *Q.J.E.* 119 (2): 403–56.
- Jacobs, James B., Coleen Friel, and Robert Raddick. 2001. *Gotham Unbound: How New York City Was Liberated from the Grip of Organized Crime*. New York: NYU Press.
- Jin, Ginger Zhe, and Phillip Leslie. 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards." *Q.J.E.* 118 (2): 409–51.
- Keenan, Jillian. 2014. "A Better Way to Take out the Garbage?" *New Yorker*, January 9. <http://www.newyorker.com/business/currency/a-better-way-to-take-out-the-garbage>.

- Kelly, Robert J. 1999. "Dirty Dollars: The Carting and Waste Disposal Industries." In *The Upperworld and the Underworld*, edited by Robert J. Kelly, 135–55. New York: Springer.
- Krasnokutskaya, Elena. 2011. "Identification and Estimation of Auction Models with Unobserved Heterogeneity." *Rev. Econ. Studies* 78 (1): 293–327.
- Krasnokutskaya, Elena, and Katja Seim. 2011. "Bid Preference Programs and Participation in Highway Procurement Auctions." *A.E.R.* 101 (6): 2653–86.
- Krishna, Vijay. 2009. *Auction Theory*. Oxford: Elsevier.
- Lewis, Gregory, and Patrick Bajari. 2011. "Procurement Contracting with Time Incentives: Theory and Evidence." *Q.J.E.* 126 (3): 1173–211.
- Li, Tong, and Quang Vuong. 1998. "Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators." *J. Multivariate Analysis* 65 (2): 139–65.
- Li, Tong, and Xiaoyong Zheng. 2009. "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions." *Rev. Econ. Studies* 76 (4): 1397–429.
- Lizzeri, Alessandro. 1999. "Information Revelation and Certification Intermediaries." *RAND J. Econ.* 30 (2): 214–31.
- Mankiw, N. Gregory, and Michael D. Whinston. 1986. "Free Entry and Social Inefficiency." *RAND J. Econ.* 17 (1): 48–58.
- Matzkin, Rosa L. 2007. "Nonparametric Identification." In *Handbook of Econometrics*, vol. 6B, edited by James J. Heckman and Edward E. Leamer, 5307–68. Amsterdam: North-Holland.
- McCall, John Joseph. 1970. "Economics of Information and Job Search." *Q.J.E.* 84 (1): 113–26.
- Moraga-González, José Luis, Zsolt Sándor, and Matthijs R. Wildenbeest. 2014. "Prices, Product Differentiation, and Heterogeneous Search Costs." Working paper.
- Nguyen, Thuy T. T., and Bruce G. Wilson. 2010. "Fuel Consumption Estimation for Kerbside Municipal Solid Waste (MSW) Collection Activities." *Waste Management and Res.* 28 (4): 289–97.
- Rosengren, Cole. 2015. "City Weighs Reining in Private Garbage Collectors." *City Limits*, May 19. <http://citylimits.org/2015/05/19/city-weighs-reining-in-private-garbage-collectors/>.
- Rubinstein, Ariel, and Asher Wolinsky. 1987. "Middlemen." *Q.J.E.* 102 (3): 581–93.
- Rust, John, and George Hall. 2003. "Middlemen versus Market Makers: A Theory of Competitive Exchange." *J.P.E.* 111 (2): 353–403.
- Salop, Steven, and Joseph Stiglitz. 1977. "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion." *Rev. Econ. Studies* 44 (3): 493–510.
- Sorensen, Alan T. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *J.P.E.* 108 (4): 833–50.
- Spulber, Daniel F. 1996a. "Market Making by Price-Setting Firms." *Rev. Econ. Studies* 63 (4): 559–80.
- . 1996b. "Market Microstructure and Intermediation." *J. Econ. Perspectives* 10 (3): 135–52.
- Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *A.E.R.* 79 (4): 700–712.
- Stigler, George J. 1961. "The Economics of Information." *J.P.E.* 69 (3): 213–25.
- Zax, David. 2014. "Dividing and Conquering the Trash." *New York Times*, October 25. http://www.nytimes.com/2014/10/26/business/dividing-and-conquering-the-trash.html?gwh=3238DE6F0B98ED349E78CCC23645724D&gwt=pay&assetType=nyt_now&_r=0.