Financial Intermediation and the Costs of Trading in an Opaque Market

Richard C. Green
Tepper School of Business, Carnegie Mellon University

Burton Hollifield
Tepper School of Business, Carnegie Mellon University

Norman Schürhoff
HEC—University of Lausanne and Swiss Finance Institute

Municipal bonds trade in opaque, decentralized broker-dealer markets in which price information is costly to gather. We analyze a database of trades between broker-dealers and customers in municipal bonds. These data were only released to the public with a lag; the market was opaque. Dealers earn lower average markups on larger trades, even though dealers bear a higher risk of losses with larger trades. We estimate a bargaining model and compute measures of dealer's bargaining power. Dealers exercise substantial market power. Our measures of market power decrease in trade size and increase in the complexity of the trade for the dealer. (*JEL* G0, G24)

Regulators, researchers, and practitioners have long debated the costs and benefits of centralization and transparency in financial markets. The fragmentation and lack of transparency may create opportunities for intermediaries to develop and exploit local monopoly power. Centralization may avoid such monopoly power locally, but it may also stifle innovation and entrench intermediaries.

Analysts have not yet resolved these questions because dealers generally resist publication of their transactions in fragmented and opaque markets. Most empirical studies of the costs of trading have been in settings where trading is relatively centralized and price information is readily available to buyers and sellers.

We analyze a database that allows us to observe the trading of broker-dealer intermediaries in the municipal bond market. We develop a simple theoretical model that decomposes the dealer's gains on a trade into the

We thank Tal Heppenstall and John Roll of PriMuni LLC for providing the data set and providing information about the structure of the municipal market. We thank participants at a number of seminars and conferences, an anonymous referee, the editor (Joel Hasbrouck), Bruno Biais, David Brown, Patrick Cusatis, Ron Goettler, Clifton Green, Larry Harris, Marc Lipson, Kai Li, Michael Piwowar, Erik Sirri, and Chester Spatt for useful comments on earlier drafts. Address correspondence to Burton Hollifield, Tepper School of Business, Carnegie Mellon University, Schenley Park, Pittsburgh, PA 15213-3890, or e-mail: burtonh@andrew.cmu.edu.

doi:10.1093/rfs/hhl012 Advance Access publication June 30, 2006

cost of facilitating the trade, a zero-mean forecast error, and a measure of the dealer's bargaining advantage or market power. Our estimates of the structural model parameters produce measures of dealers' market power and show how it varies with the type of trade. Our analysis thus contributes to a better understanding of the consequences of fragmentation and transparency by measuring the risks and rewards earned by intermediaries in an opaque and fragmented market of the sort that has so far been the most difficult to study.

The market for municipal bonds in the United States is a large and important financial market by any measure. Aggregate municipal bond holdings currently amount to roughly \$2 trillion.[1] New issues in recent years have exceeded \$400 billion per year. Along with its size, the municipal market has been remarkable for its fragmentary nature and lack of transparency. These characteristics are attributable to several factors.

First, there are many relatively small issuers of municipal bonds. In recent years there have been between 10,000 and 15,000 separate municipal bond issues per year. In 2003 the average issue size was \$26 million.[2] Municipal bonds are typically issued in series, and larger issues, in turn, are likely to consist of 10–30 separate bonds with different maturities. Second, because municipals are tax-exempt, they are in large measure held by individual retail investors (33% in 2004), by insurance companies (12%), or in bank and personal trusts (5%). Municipal bonds are therefore unlikely to trade frequently.

Third, municipals trade in decentralized broker-dealer markets. There is no centralized exchange providing pre-trade transparency. To obtain quotes, a buyer or seller must call multiple dealers or solicit bids from them. Comparison shopping is relatively costly. Fourth, many states only exempt their own municipalities' bonds from state income tax, further contributing to local segmentation. Finally, for political reasons, or because of local cost advantages, smaller regional firms underwrite many municipal issues. The lack of a large intermediary with access to many different bonds may further increase the costs of efficiently matching buyers and sellers.

The size of spreads in the municipal market has attracted the attention of regulators, the press, and the investing public in the last few years. Many argue that the spreads are unreasonably high. A popular web site (http://www.MunicipalBonds.com) lists the "Worst Ten Spreads" for trades in various categories, along with "Red Flags" and "Worst Spread Reports" based on Municipal Securities Rulemaking Board (MSRB) reports. MSRB rules require that trades with customers should only be

[1] The source of these aggregate statistics on the municipal market is The Bond Market Association (http://www.bondmarkets.com).

[2] Data on the number of underwritings are available at http://www.financialservicesfacts.org.

executed at prices that are "fair and reasonable, taking into consideration all relevant factors." Both National Association of Securities Dealers (NASD) and the Securities and Exchange Commission (SEC) have examined specific cases and complaints concerning municipal bond trades. The difficulty for researchers and policy makers lies in identifying "all relevant factors" and quantifying when variation attributable to those factors is "fair and reasonable."

The SEC ruled that markups on trades ranging from 1.42 to 5% "substantially exceeded accepted industry practice," whereas NASD stated that transactions with spreads between 7 and 27% were not "fair and reasonable."[3] Such blanket rules, because they ignore the variation in the costs of financial intermediation, will necessarily be conservative and identify only extreme cases. Our structural model allows us to disentangle competitive compensation for the costs of intermediating trades from the exercise of market power.

Our sample comes from the MSRB who gathered the sample as a first step in their efforts to move municipal bond trading toward a more transparent system. The sample records every transaction in municipal bonds by registered broker-dealers between May 1, 2000, and January 10, 2004. The sample contains a large number of transactions—over 26 million trades—covering a longer time than in a typical microstructure, transaction-based sample.

These data are only made public after a time lag. Initially the lag was one day for bonds that traded more than four times during the day and one month for all other trades. Over the sample period the time lag has been shortened, and the scope of the daily reports has been increased in several steps.[4]

The sample does not identify the specific broker-dealer associated with a given trade but does record whether the transaction was between dealers, or between a dealer and a customer. We can measure the markups implicit in transaction prices between dealers, in aggregate, and their customers and study the determinants of the markups, such as the price concessions to quantity.

Because we cannot directly link individual trades to specific dealers, we must infer their trades and profits indirectly. For example, we study purchases that can be matched with subsequent sales using a first-in-first-out (FIFO) rule. More narrow filters, such as using pairs of trades where the same par value of the same bond is purchased from a customer and then sold to a customer in a very short time, provide more confidence that we are identifying two sides of the same transaction, at the expense of greater selection bias and

[3] See "Regulators Know What Bond Dealers Did Last Summer," by Joe Mysak, *Bloomberg News Service*, September 3, 2003, "Munis: What's a Fair Price?" by Dean Foust, *Business Week*, July 7, 2003, and SEC "Opinion of the Commission," Admin. Proc. File No. 3–9499.

[4] Since the sample period ended, the MSRB reports have started to provide close to real-time post-trade transparency. Transactions are available on http://www.investingbonds.com.

less representative measures of profits to dealers in aggregate. Our empirical findings are robust to the method of matching trades.

The measures we obtain of profits to dealers correspond to round-trip transaction costs to customers. Their distribution is highly skewed, and much more variable for smaller trades. Dealers lose money with much greater frequency on large, institutionally sized trades.

Descriptive empirical methods alone are of limited use in identifying that portion of dealers' profits coming from market power. We therefore develop a simple theoretical model of the interaction between dealers and their customers in which the expected profits to the dealer reflect both the dealer's costs and his bargaining power relative to the customer. Both of these, in turn, can be parameterized as functions of observable variables and estimated as a stochastic frontier model. The dealer's cost is the frontier, which represents the expected markup the customer would obtain if dealers were always driven to their reservation values, as they would be if the provision of dealer services were perfectly competitive. The observed markup differs from dealer's costs by a zero-mean forecast error and a one-sided error. The one-sided error reflects the distribution of sellers' reservation values and dealer's bargaining power.

Our structural model provides the decomposition of round-trip trading costs into dealer's costs and dealer's market power. The identification of the model depends on our assumptions about the functional forms in the structural model. The model's success in fitting the data is because of a high degree of conditional skewness in the distribution of markups dealers earn. The skewness is much more dramatic for smaller trades, suggesting dealers can exercise greater market power in their dealing with such customers.

Our empirical results suggest that dealer's costs depend on liquidity measures for the bond and market, the size of the trade, interest rate conditions, and expectations about how the trade will be processed. The portion of average profits attributed by the estimates to the dealers' advantage in bargaining power is substantially higher than the portion attributed to the actual cost of intermediating the trades. The dealers' market power is highest for small- to medium-sized transactions. Such transactions are presumably with less sophisticated retail investors: a finding consistent with the theoretical arguments in Duffie, Gârleanu, and Pedersen (2005) that less sophisticated investors face higher markups in a bargaining market. The market power of dealers as a group is also higher in situations with a higher probability that more extensive intermediation will be required, such as moving the trade through other dealers, or breaking it into smaller blocks.

The lack of centralized recording and the preservation of price information have limited the study of trading costs in bond markets. Exceptions are studies by Hong and Warga (2000), Schultz (2001), and

Chakravarty and Sarkar (1999). These authors all studied a sample of trades executed on behalf of large insurance companies. Their data identified the institutions and dealers participating in the trade and whether it was a buy or a sell. The authors used various techniques to estimate the bid–ask spread implicit in the observed trades and then related the measured trading costs to the characteristics of the bonds and the current market conditions. Of these studies, only Chakravarty and Sarkar (1999) included municipal bonds in their analysis.

Our sample differs from that used in the previous studies in several respects. First, the sample is comprehensive. It includes all trades by registered broker-dealers in municipal securities, rather than only the trades initiated by a subset of institutional investors. We have data on over 26 million transactions in over 1 million bonds issued by over 46,000 entities. Of these transactions, 5.3 million are purchases from customers in seasoned bonds, and we are able to match 4.5 million of the transactions with subsequent sales through our broadest matching rules.

Hong and Warga (2004) also used the MSRB data over a much shorter sample period and estimated effective spreads by averaging differences between prices on purchases from and sales to customers in the same bond on the same day. They also found spreads exceed 2% on smaller, retail-sized trades.

The article closest to our work is that by Harris and Piwowar (2005). They estimated trading costs for municipal bonds using the MSRB data, though over a shorter time period. Their approach is very different from ours, reflecting differences in goals. The two articles share an interest in measuring how large the costs of trading in this market are and in understanding how these costs vary with the size and other characteristics of the trade. The articles also share a central difficulty: although there are many trades in many bonds, trading in any one bond is typically very infrequent.

Harris and Piwowar (2006) estimated a time-series model of the trading cost for each bond with a minimal number of trades during the sample period, augmented with a factor model that uses bond market indices to infer movements in the intrinsic value of the bonds between observed trades, which may be months apart. They then analyzed how the estimated costs vary cross-sectionally for different types of bonds. We focus on trades that can reasonably be assumed to represent two sides of a single intermediated transaction and employ a structural model to decompose the cost faced by a customer into a portion that represents the cost the dealer incurs and a portion attributable to the dealer's market power.

These very different approaches produce roughly similar estimates of average costs of trading, and both make clear that smaller trades are much more costly than large ones. The Harris and Piwowar's (2006) results provided a richer description of how the costs vary with bond characteristics. That article is purely empirical in its goals, however. We use a theoretical model to seek evidence that the high costs of trading are due to dealer's market power and to find out how the exercise of market power depends on the characteristics of the trade. The presence of market power and the way dealers use it are central policy and research concerns in markets that lack transparency.

The remainder of this article is organized as follows. Section 1 describes the institutional setting and our sample. In Section 2 we study dealer's profits on round-trip transactions and illustrate their dependence on size and other characteristics of the trade. In Section 3 we present a structural model that allows us to decompose the profits into portions attributable to dealer's costs and market power. We also discuss in detail how the model is identified. We report our empirical estimates of the structural model in Section 4. In Section 5 we incorporate information on how trades are processed. The final section concludes. The appendixes contain detailed descriptions of procedures we used to filter the data, variable definitions, and a discussion of model identification.

1. The Municipal Bond Market and the Sample

We have a sample from the MSRB reporting all trades carried out by broker-dealers from May 1, 2000, through January 10, 2004. The sample was gathered as part of the MSRB's efforts to improve transparency in the municipal bond market and will eventually be incorporated into a system that makes transaction prices publicly available to market participants. The sample is self-reported, and there are many systematic data errors, especially early in the sample period when member firms were still becoming accustomed to the system. In Appendix A we describe in detail the procedures we use to identify and filter out data errors.

For each trade, we observe the price, the date and time of the trade, the par value traded, and whether the trade was a dealer buying from a customer, a dealer selling to a customer, or a trade between dealers. The yield at which the trade took place is sometimes reported. Information is also provided about the traded bond including the CUSIP number, a text description of the bond, the date interest begins accruing (called the "dated date"), the coupon rate, and the bond's maturity. Often some of the fields are empty, but by searching other trades for the CUSIP number, we can generally determine the missing information.

In addition to the information in the MSRB data set, we have obtained additional descriptive information about the individual bond issues through NASD and the web site "investingbonds.com." The information includes the bond rating and call schedule.

Table 1 provides descriptive statistics on the transactions in the MSRB database. This Table 1 includes all transactions in the data and describes

Table 1
Descriptive statistics

	All transactions	Sales to customers	Purchases from customers	Transactions between dealers
		All issues		
Observations (million)	26.80	15.68	5.69	5.44
Issues (million)	1.08	1.06	0.71	0.64
Issuers (thousands)	46.50	46.13	43.09	37.05
Total value (trillion)	11.26	5.79	3.92	1.55
Average par (million)	0.42	0.37	0.69	0.29
Coupon (%)	5.12	5.03	5.37	5.14
Age (years)	3.78	3.06	5.62	3.92
Maturity (years)	14.96	15.85	13.85	13.55
		New issues		
Observations (million)	7.93	5.84	0.38	1.71
Issues (million)	0.61	0.60	0.14	0.33
Issuers (thousands)	26.09	25.88	17.18	18.61
Total value (trillion)	3.70	2.31	0.61	0.78
Average par (million)	0.47	0.40	1.64	0.46
Coupon (%)	4.59	4.62	4.66	4.47
Age (years)	0.03	0.03	0.05	0.03
Maturity (years)	15.27	15.80	14.36	13.64
		Seasoned issues		
Observations (million)	18.88	9.84	5.31	3.72
Issues (million)	0.65	0.63	0.63	0.39
Issuers (thousands)	41.71	41.27	41.29	34.35
Total value (trillion)	7.56	3.48	3.31	0.77
Average par (million)	0.40	0.36	0.63	0.21
Coupon (%)	5.36	5.29	5.42	5.46
Age (years)	5.33	4.83	6.00	5.70
Maturity (years)	14.83	15.88	13.81	13.50

This table reports descriptive statistics for the MSRB transaction data from May 1, 2000, to January 10, 2004. The total value is the market value of the trades over the sample period. Average par is the average par value of the bonds in a transaction. Coupon, age, and maturity are sample averages.

trades occurring in the first 90 days after the dated date for bonds that are issued during our sample period.[5] We refer to such trades as "new issues."

Sales to customers exceed buys from customers, primarily because of activity in new issues. Syndicates of broker-dealers purchase new bond issues, and the resulting inventory is distributed to customers primarily during the first 90 days. The dollar value of trades in newly issued bonds is \$3.7 trillion, only 32% of the total value of all trades. The difference between the value of sales to customers and buys from customers for newly issued bonds, \$1.7 trillion, is 90% of the corresponding difference

[5] For its monthly trade summaries, the MSRB defines "new issue trades" as "trades where the difference between the trade date and the dated date is less than or equal to 4 weeks (28 days)." We chose a longer time horizon, because our data suggest that transaction volume has not yet flattened out one month after issuance. Substantial transaction volume in new issues through the when-issued market starts as early as 20 days before the dated date.

for all trades, even though transactions in new issues are less than a third of all transactions.

The value of sales to customers is approximately equal to buys from customers in seasoned issues. The broker-dealers tend to buy bonds from customers in larger par amounts and then sell the bonds in smaller blocks. For new issues, the average purchase from a customer is close to four times larger than the average sale to a customer, and for seasoned issues, the average purchase from a customer is a bit less than three times the average sale to a customer. Interdealer trades are of intermediate size.

Municipal bonds are actively traded in a "when issued" market and also actively traded immediately after they are issued. Once the bonds find their way into retail and mutual fund portfolios, the volume of trade drops off dramatically. The MSRB reports that from March 1998 to May 1999, 71% of the outstanding issues did not trade at all. We compute the empirical probability of trade as a function of the age of the bond. In the first month of issuance, all the bonds trade, in the second month approximately 15% of the bonds trade and from two to six months from issuance, less than 10% of the bonds in our sample trade at all. The probability then rises somewhat so that by four years from issuance, roughly 15% of the bonds in the sample trade at least once during a month.

We also compute the average number of trades per day, conditional on the bond trading on that day. After the bond is reasonably seasoned, the average number of trades stays close to two, suggesting that trades through broker-dealers in a given bond are typically limited to the intermediation of a specific transaction.

Our empirical tests in subsequent sections focus on liquidity provision for seasoned bonds—bonds more than 90 days from issuance—because volume has typically stabilized for seasoned bonds. The value of purchases from customers roughly equals sales for seasoned bonds. Because we see both sides of most transactions in seasoned bonds, we can reliably measure the profitability of the flows to the dealers for transactions in seasoned bonds. Our results are, therefore, most informative about the typical experience of an individual or institution holding bonds who wants to sell them and thus demands liquidity.

Table 2 lists the independent variables we employ, with the details of their sources and construction. These include the characteristics of the transaction itself, such as the size and whether the initial purchase was sold by splitting it into smaller blocks, attributes of the bond, such as rating and seasoning, and general market conditions, such as the level of interest rates and volume in bonds from the state in question.

In Table 3 we list the attributes of the bonds in our sample and the percentages of the bonds and of the transactions of each type. Although our data sources provide ratings for less than 10% of the bonds, bonds with ratings amount to over 40% of the transactions. The vast majority of

Table 2
Definitions of the explanatory variables

Variable	Description
Markup	Difference between the purchase and the resale price divided by the purchase price, truncated at 0.5 and 99.5%
Transaction frequency	The average number of purchase transactions per day over the minimum of the bond's age, the start of the sample, or the last 365 days, ending the day prior to the transaction
State volume	The total value of trades per day in bonds from that state over the last seven days ending the day prior to the transaction. Measured in billions of dollars
State order imbalance	Purchases from customers in seasoned bonds issued by municipalities in that state less sales to customers per day over the last seven days ending the day prior to the transaction. Measured in billions of dollars
State new issues	Sales of new bonds in the same state per day over the last seven days ending the day prior to the transaction. Measured in billions of dollars
Treasury yield	Closing yield of the maturity matched treasury bond on the day prior to the purchase transaction. Measured in percent
Treasury volatility	Parkinson (1980) estimate of volatility for the maturity-matched treasury yield over a thirty day window ending the day prior to the transaction. Measured in basis points
	Dummy variables
Coupon zero or missing	Indicator variable for an original issue pure discount bond, or if coupon information is missing after search over database and imputation
No splitting, interdealer	Bonds are sold in a block through at least one other dealer
Splitting, no interdealer	Bonds are sold to customers in multiple blocks with no interdealer trades in the meantime
Splitting and interdealer	Bonds are sold to customers in multiple blocks with interdealer trades in the meantime
Rating available	Bond issue is rated and the rating is available to us
Rating A-BBB	Bond rating is between A+ and BBB-
Rating BB-B	Bond is speculative grade. Bond rating is between BB+ and B-
Rating CCC-D	Bond has junk or default status. Bond rating is between CCC+ and D-
Callable bond	Bond is callable between the transaction date and maturity
Taxable bond	Bond is not tax exempt
Pre-refunded bond	Bond has been pre-refunded. These bonds are backed by trusts invested in treasuries, and are riskless
Guaranteed bond	Bond is guaranteed. Bonds backed by an entity with credit risk; the bonds have credit risk
Insured bond	Bond has insurance from a municipal bond insurance agency, or from a specific entity
Letter or line of credit	The bond is backed by a letter or line of credit
Revenue bond	Revenue bond
Tax revenue bond	Bond backed directly by tax revenues from a specific source
General obligation bond	General obligation bond

Table 2
(continued)

Variable	Description
Certificates of participation	Certificate of Participation
County issuer	Issuer is a county
School district issuer	Issuer is a school district
Development authority	Issuer is a development authority
Financial authority	Issuer is a financial authority
Housing authority	Issuer is a housing authority
Sewer authority	Issuer is a sewer authority
Redevelopment	Issuer is a redevelopment authority
Utilities construction	Bond issued to finance utilities construction
Improvement	Bond issued for facilities improvements
Facilities construction	Bond issued to finance facilities construction
Refinancing	Bond used to refinance another issue

This table describes the conditioning variables.

Table 3
Bond characteristics

	Bonds (%)	Full sample (%)
Rating available	8.7	42.4
Rating AAA	73.2	74.1
Rating AA	18.7	14.5
Rating A-BBB	7.4	9.6
Rating BB-B	0.5	1.2
Rating CCC-D	0.2	0.6
Callable bond	31.7	45.7
Taxable bond	1.9	2.4
Pre-refunded	0.8	2.3
Guaranteed bond	0.6	0.6
Insured bond	44.2	49.8
Letter or line of credit	1.1	4.2
Revenue bond	36.1	59.6
Tax revenue bond	5.3	4.4
General obligation bond	2.1	2.2
Certificates of participation	7.9	3.9
County issuer	22.6	21.3
School district issuer	21.0	8.8
Development authority	3.0	4.6
Finance authority	2.1	3.5
Housing authority	0.8	0.6
Sewer authority	4.4	4.0
Redevelopment	1.4	1.3
Utilities construction	3.3	2.7
Facilities improvement	8.3	5.8
Facilities construction	6.8	13.2
Refinancing	24.8	29.8

This table describes the bonds' characteristics. The first column lists the percentage of the bonds appearing in our sample falling into the category in question. The second column gives the percentage of transactions in the full sample. The percentages for the different ratings categories are conditional on a rating being available.

them are rated AAA. Over 40% of the bonds in our sample are insured and therefore likely to be rated AAA, and almost 50% of the transactions are in insured bonds. This illustrates that many of the bonds for which we do not have rating information are in fact rated. It also points to an important feature of the municipal market generally, which makes the size and heterogeneity of the transaction costs we see that much more puzzling. The insured bonds are virtually perfect substitutes with each other from the standpoint of credit risk.

2. The Magnitude and Characteristics of Dealer Markups

Broker-dealers in the market for seasoned municipal bonds typically supply liquidity by purchasing blocks of bonds from customers and then reselling the bonds, either in a block or split into smaller trades. From May 1, 2000, to January 10, 2004, a total of 5,313,692 purchases from customers occurred in seasoned municipal bond issues.

Our data do not identify matched buy–sell transactions for specific bonds, nor the specific dealers carrying out those transactions. We observe bonds flowing into and out of the hands of dealers as a group. In many cases, however, the sequence of transactions in a given bond issue leaves little doubt that they are two ends of a transaction intermediated by a specific dealer. We focus on three methods for measuring the profits to dealers from the transactions. The narrowest measure isolates transactions that are most obviously two sides of the same transaction. The broader measures capture more trades and thus can be viewed as more representative of the profits dealers earn as a group, but at a cost of less reliability in identification.

First, we study trades where a buy from a customer is followed by a sale to a customer in the same bond for the same par amount on the same day with no intervening trades in that bond. We refer to such trades as "immediate matches." For immediate matches, no added inventory is carried on the books of the dealers at the end of the day. As shown in the first row of Table 4, there are roughly 750,000 such matched pairs in the data set.

Of course, some of the immediate matches may be situations where the initial purchase is by one intermediary, the subsequent sale is by another, and the par value of the trades is identical because of a coincidence. The information from the previous section about the frequency of trade, however, suggests that such events are quite unlikely, particularly when the time between the two trades is short. The mean number of trades per day, conditional on a bond trading once, is close to two. If, in a short period of time, we know one of these trades is a buy, the other is a sale, and they both involve the same par value, it is very likely that they are two sides of a trade intermediated by a single dealer.

Table 4
Sample selection: Descriptive statistics

	Immediate matches	Round-trip transactions	FIFO sample	Full sample	All
Observations (million)	0.75	1.96	1.87	4.58	5.30
Gross markup (%)	1.33	2.03	1.18	1.69	–
Net markup (%)	1.33	2.00	1.11	1.61	–
Yield: purchase (%)	4.84	4.76	5.00	4.86	4.84
Yield: sale (%)	4.10	4.10	4.40	4.21	–
Par (thousands)	25.00	25.00	60.00	35.00	35.00
Maturity (years)	11.24	9.91	17.65	13.05	12.82
Coupon (%)	5.13	5.00	5.00	5.00	5.00
Age (years)	6.48	5.20	4.50	5.10	5.15
Transaction frequency (1/day)	0.03	0.02	0.12	0.04	0.04
State volume (billion)	0.20	0.21	0.22	0.21	0.21
State new issues (billion)	0.06	0.06	0.06	0.06	0.06
Treasury yield (%)	4.72	4.69	5.05	4.86	4.81

This table reports descriptive statistics on the three subsamples, the full matched sample, and all data remaining after applying the various data filters. The numbers shown are the median values over the corresponding sample. Immediate matches are those dealer purchases of bonds followed by a sale of the same par amount on the same day with no intermediate transactions. Round-trip transactions include all purchases that can be matched with subsequent sales, with no intermediate purchases, excluding immediate matches. The FIFO sample includes all remaining purchase transactions in seasoned bonds that can be matched with subsequent sales using a first-in-first-out rule. The full sample consists of the immediate matches, the round-trip transactions, and the FIFO sample. The last column in the body reports statistics for all purchase transactions that remain after applying the various data filters including those that did not get matched. Gross markup is the difference between the purchase price and sales price, as a percentage of the purchase price. Net markup is the gross markup less the return on a maturity-matched municipal index.

Second, we consider a broader set of transactions where buys from customers are followed by transactions that clear dealer inventories in aggregate. In each case, a buy from a customer is followed by one or more sales to customers equal in par value to the initial purchase with no intervening purchases from customers. The purchases and sales may be spread out through time, and there may be intermediate trades between dealers in the bond. The profits earned on such transactions represent rewards to the dealers as a group for the intermediation they provide. We refer to such transactions as "round-trip transactions." As Table 4 summarizes, there are almost 2 million round-trip transactions that are not immediate matches.

The most inclusive selection procedure we use, or the "FIFO sample," consists of all the purchases from customers that over our sample period could be matched with sales to customers using a FIFO rule and were not already selected by one of the narrower criteria. For the FIFO sample, there may be intervening purchases from customers, but unless those purchases are selected by the round-trip criteria, the first sales are assigned to the bonds that came into dealer inventories first. The FIFO sample, then, is most likely to assign sales to purchases made by another dealer but will be more representative of profits to the dealers as a group. The FIFO selection rule identifies an additional 1.9 million transactions. The "Full sample" combines the three sets of transactions and covers 86% of the 5.3 million purchases from customers in our database. The remainder were purchases from customers that remained in inventory at the end of the sample period, or where the offsetting sale was not recorded because of a clerical mistake. Appendix B provides more details on the matching procedures used to construct the samples.

We use the full sample in most of the empirical results we report, but we provide many robustness checks to illustrate that the broad conclusions would follow from analyzing any of the subsamples. It is important to recognize that the different selection criteria involve endogenous choices. Indeed, any transaction-based sample involves a decision to trade. Our purpose in using different measures of dealer's profits, however, is not to draw distinctions between their determinants, but rather to establish that our results and conclusions are reasonably robust to how the trades are selected.

Table 4 summarizes the four sets of trades and reports the medians of the measured profits on the trades and some of the explanatory variables associated with the trades that we use in subsequent analyses. The "Gross markup" is the dealer gain or loss on the trades as a percentage of the purchase price. The "Net markup" adjusts the markup for yield curve movements. For this purpose, we have obtained data for the Lehman Brothers Municipal Bond Index for various maturities (1, 3, 5, 7, 10, 15, 20, and 22+ years) from Datastream. We use a linear spline to interpolate values for other maturities. To arrive at the net holding-period profits, we subtract the maturity-matched price change of the corresponding municipal bond index from the gross markup. We truncate the resulting distribution at the 0.5 and 99.5% quantiles, because there are some very extreme observations that appear to have been the result of clerical errors when the data were entered.

The markups measure the cost of round-trip transactions to an investor, or the profits to a dealer from buying and disposing of the bonds. For an investor holding the bond and seeking liquidity by selling it, one side of the round-trip cost is already sunk, and the relevant marginal transaction cost is half the markup we measure.

Table 4 summarizes that median markups are between 1.3 and 2%. The table also summarizes that median yields on municipal bonds are less than 5% in our sample—transferring a bond from one investor to another, therefore, involves surrendering much of a year's return to the intermediary.[6]

The different selection criteria do deliver trades with different characteristics. Bonds that trade more frequently are more likely to have

[6] The comparison ignores tax differences for retail investors between the capital loss implicit in the cost of selling the bonds and the tax exempt coupon income, and so may slightly overstate the relative cost.

intervening purchases before the position is sold off, particularly if the initial purchase is not immediately disposed of. Given the lower liquidity, we do not find it surprising that the measured markups are also larger for the round-trip transactions. The two stricter criteria—the immediate matches and round-trip transactions—select trades involving lower par amounts, and bonds where daily volume in bonds from the same state was lower.

Table 5 indicates that estimates of the time before bonds purchased in a block clear dealer inventories for the two broader sample criteria: the round-trip transactions and the FIFO sample. Over half of the purchases from customers are sold within five business days. But the dealers as a group do appear to bear substantial risk of ending up holding bonds for long periods when they supply liquidity. The extent to which holding bonds for long periods is costly for the dealers is unclear, however, because the bonds pay competitive returns.[7]

Tables 6–8 provide information from the full sample about the markups the dealers earn, the frequency of losses, and how trades are processed. We only report this information for the full sample, but similar patterns are very evident for the narrower selection criteria and for shorter time periods. Trades are categorized by whether the initial purchase is sold off in smaller blocks ("Split" versus "No Split") and by

Table 5
Sample selection: Distribution of the time in dealer inventory (%)

Day	Round-trip transactions	FIFO sample	Full sample	All
1	40.3	11.1	38.1	33.2
2	60.1	17.3	49.1	42.8
3	70.2	22.2	55.4	48.4
5	81.4	30.3	63.5	55.5
30	98.9	64.9	85.2	74.9
360	100.0	95.4	98.1	88.6

This table reports the distribution of the time in dealer inventory obtained from the Kaplan—Meier nonparametric estimator. Round-trip transactions include all purchases that can be matched with subsequent sales, with no intermediate purchases, excluding immediate matches. The FIFO sample includes all remaining purchase transactions in seasoned bonds that can be matched with subsequent sales using a first-in-first-out rule. The full sample consists of the immediate matches, the round-trip transactions, and the FIFO sample. The last column in the body reports statistics for all purchase transactions that remain after applying the various data filters including those that did not get matched. Gross markup is the difference between the purchase price and sales price, as a percentage of the purchase price. Net markup is the gross markup less the return on a maturity-matched municipal index.

[7] For the dealers, the coupon income in municipal bonds is tax-exempt, and capital gains are taxed on an accrual basis. At the short-end of the term structure, municipal bonds have similar after-tax yields to treasuries, and at longer maturities, municipal bonds tend to have higher after-tax yields than treasuries. See, for example, Green (1993).

Table 6
Descriptive statistics for the full sample: Frequency (%)

	Full sample	*Par* ∈ (0,100*k*)	*Par* ∈ [100*k*,500*k*)	*Par* ∈ [500*k*,∞)
Number of observations (million)	4.58	3.10	0.79	0.69
No Split × No Interdealer	55.3	61.4	41.0	44.3
Split × No Interdealer	18.3	14.0	22.7	32.7
No Split × Interdealer	13.0	15.6	11.0	3.5
Split × Interdealer	13.4	9.0	25.4	19.4

This table reports the frequency distribution of the dealers' decisions to split the order and to trade with another dealer. Trades where the initial purchase is sold in more than one block are categorized as split. Trades that involve interdealer trade have trades between dealers prior to the sales that clear aggregate dealer inventory of the par amount initially purchased.

Table 7
Descriptive statistics for the full sample: Net markups–measured (%)

		Full sample	*Par* ∈ (0,100*k*)	*Par* ∈ [100*k*,500*k*)	*Par* ∈ [500*k*,∞)
Full sample	Median	1.61	2.13	0.81	0.00
	Mean	1.77	2.30	1.10	0.16
	SD	1.99	1.88	1.80	1.51
No Split × No Interdealer	Median	1.51	2.01	0.22	0.00
	Mean	1.64	2.13	0.41	−0.09
	SD	1.78	1.65	1.42	1.11
Split × No Interdealer	Median	1.51	2.30	1.41	0.00
	Mean	1.52	2.36	1.51	−0.10
	SD	1.85	1.55	1.60	1.47
No Split × Interdealer	Median	2.08	2.46	0.85	0.28
	Mean	2.42	2.75	1.14	0.44
	SD	2.48	2.49	1.87	1.68
Split × Interdealer	Median	1.77	2.41	1.56	0.91
	Mean	2.02	2.59	1.83	1.13
	SD	2.24	2.36	2.06	1.87

This table reports the associated markups. Trades where the initial purchase is sold in more than one block are categorized as split. Trades that involve interdealer trade have trades between dealers prior to the sales that clear aggregate dealer inventory of the par amount initially purchased.

whether the bonds are traded between dealers before being sold off to customers ("Interdealer" versus "No Interdealer"). For larger trades of over $100,000 par value, just over 45% are resold in a single block. For retail-sized transactions of less than $100,000 in par value, over 75% are resold in a single block. Multiple dealers are most likely to be involved in trades of intermediate size. For the full sample, roughly a quarter of the trades involve interdealer transactions, whereas for intermediate sizes this amounts to 35% of the trades.

Table 8
Descriptive statistics for the full sample: Frequency of net loss (%)

	Full sample	$Par \in (0,100k)$	$Par \in [100k,500k)$	$Par \in [500k,\infty)$
Full sample	13.5	5.7	19.3	42.2
No Split ×No Interdealer	12.8	4.8	29.1	45.8
Split × No Interdealer	18.4	3.9	11.5	51.8
No Split × Interdealer	10.7	8.7	16.4	29.9
Split × Interdealer	12.7	9.7	11.9	20.2

This table reports the frequency of net loss. Trades where the initial purchase is sold in more than one block are categorized as split. Trades that involve interdealer trade have trades between dealers prior to the sales that clear aggregate dealer inventory of the par amount initially purchased.

Table 5 reports statistics describing the percentage markup for trades of different sizes and processed in different ways. Columns in the table divide the trades according to the par value of the initial purchase in millions, rows by method of disposition. Regardless of how the trade is processed, average markups fall as the size of the trade increases, as does the standard deviation of the realized markups. Markups are negative or zero for institutionally sized trades of over $500,000, if the transactions do not involve multiple dealers. Harris and Piwowar (2006) also found the largest trades earn the dealers negative profits, using very different methods. The negative profits may indicate that the dealers are demanding liquidity rather than offering it. This would be the case if the dealers needed to obtain inventory from institutions to fill orders for retail customers. The finding may also suggest cross-subsidization, because, in equilibrium, dealers must be covering the costs and overhead of their municipal bond trading operations elsewhere, either in trades with other customer types or through other services to the same customers.

Within any size category, profits to dealers as a group are higher for trades that require more extensive or elaborate intermediation in the sense that many counterparties are involved in selling off the bonds. Trades that involve a larger number of counterparties are associated with more uncertainty than trades that involve a fewer counterparties—dealers' profits generally have higher standard deviations if the initial purchase is split or leads to intermediate trades. The causality between profits, uncertainty, and number of counterparties is ambiguous, because the method of selling off the bonds is an endogenous choice of the dealers.

For example, dealers may involve more counterparties when they anticipate more difficulties clearing their inventory ex-ante, or they may involve more counterparties ex-post, when market conditions have turned against them and they become more desperate to sell. The positive association between profitability and both uncertainty and the number of counterparties suggests that dealers involve more counterparties when they anticipate more difficulties clearing their inventories ex-ante. Yet another possible explanation for higher markups when multiple dealers or counterparties are involved is that, for the dealers as a group, it circumvents norms ruling out "unreasonable" markups. Violations of these norms would be easier to detect if the bonds are bought and sold in blocks of the same size through a single intermediary.

Overall, the evidence in Table 7 seems consistent with dealer markups being determined by the uncertainty and costs dealers bear, rather than market power. There are likely economies of scale in handling trades, and percentage markups fall with trade size. More service is provided by the intermediaries when more counterparties are involved, and larger markups are evident for these trades. Trades with higher average markups also have higher standard deviation.

Nevertheless, Tables 6–8 also provide evidence that bargaining power is tilted toward larger trades. Although standard deviation is higher for smaller trades, dealers rarely lose money on them. Table 8 reports loss frequencies for the various categories of trades. Dealer losses are most frequent for the categories of trades on which markups are the least variable. There is a high level of skewness in the markups for all trade categories, and for large trades, the distribution is simultaneously more concentrated and centered at a lower level.

This skewness is very evident in Figure 1 showing the frequency distributions for three different trade-size categories for the full sample. The plots, which all have the same scale, are kernel density estimates. The estimates are computed using the Epanechnikov kernel with the Silverman automatic bandwidth selection method.[8] It is obvious from the plots that the higher loss frequencies for large trades are not the result of the realized percentage profits being more spread out. Their distribution is much more concentrated but centered at a much lower point.

The combination of these three characteristics—lower percentage profits, higher probability of losses, and less variable profits—suggests that the largest traders transact with the dealers on more attractive terms. The next section, which reports the estimates of parameters that summarize dealer's market power, provides some additional evidence along these lines.

Figure 1 also provides some evidence on how pervasive extremely high markups are for trades of different sizes. On relatively small transactions, a substantial portion of the distribution of realized spreads exceed 5% (the NASD's "cap" for equity trades), whereas very few

[8] Härdle (1990) provided a textbook introduction to kernel smoothing methods.

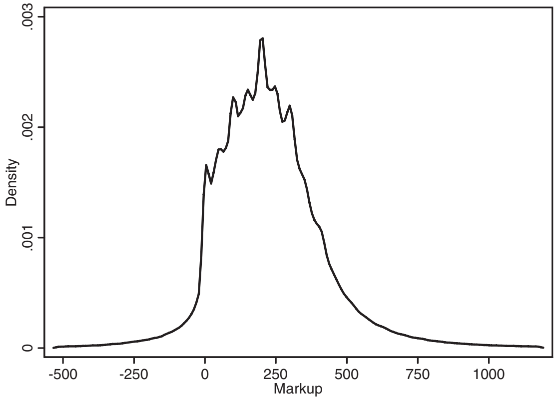

Panel A $Par \in (0, 100k)$

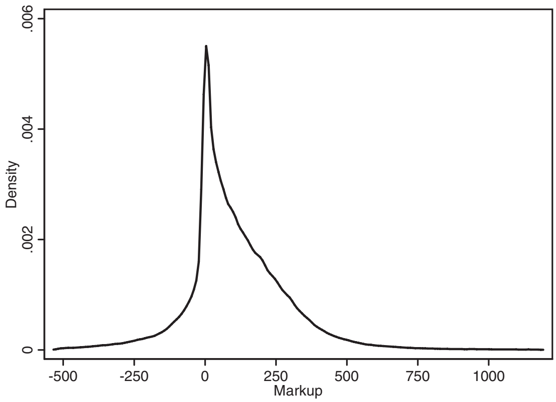

Panel B $Par \in [100k, 500k)$

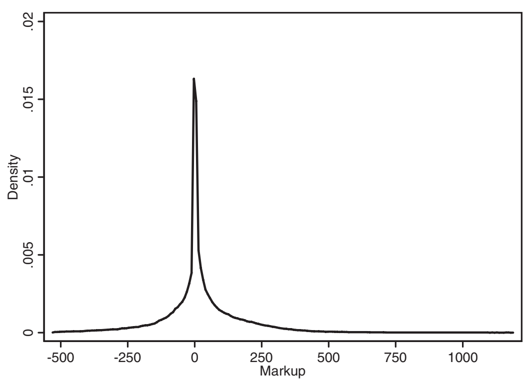

Panel C $Par \in [500k, \infty)$

Figure 1
Distribution of net markups for different transaction sizes
The distributions are estimated using Epanechnikov kernels with the Silverman automatic bandwidth selection. Net markups are measured in basis points and are equal to the difference between the purchase price and sales price, as a percentage of the purchase price, less the return on a maturity-matched municipal index.

trades occur at spreads in excess of 8%. Except for the very largest trades, there is considerable probability mass above 2%. In summary, relative to the norms of fairness and reasonableness cited in SEC opinions and NASD complaints, there are substantial numbers of trades that occur at high spreads, but the extremely high spreads in the range of 20–50% reported in the press are likely to be data errors or extremely unusual events.

3. Dealer's Market Power and Transaction Size

In the previous section we established that the markups dealers earn on round-trip transactions in municipal bonds are large. They differ in magnitude and variability across trades with different characteristics. To what extent is this evidence that dealers exercise market power in their interaction with certain classes of customers, versus evidence that the costs and risks of intermediating the trades differ with these characteristics? We speak to the question by developing and estimating a structural model that decomposes the observed markup into components associated with the cost to the dealer and the dealer's relative bargaining power with the customer.

3.1 The model

To keep the exposition simple, we will develop the model in terms of price levels and explain later how it can be expressed in percentage terms. Most round-trip transactions in seasoned municipal bonds are initiated by customers wishing to sell particular bonds. Accordingly, consider transactions between a dealer and a customer currently holding bonds she wishes to sell.

We assume that the seller arrives with a reservation value v, which represents her outside opportunities. The reservation value can be viewed, for example, as the continuation value associated with contacting a different dealer and searching for better terms of trade. The dealer's decision about whether to trade depends on her expectations about the price she will eventually obtain on selling the bond and the costs she anticipates in intermediating the trade. We denote the eventual sales price as p, and its expectation as $E(p \mid X)$, where X is a set of conditioning variables, observable to the dealer and the seller. The anticipated cost of intermediating the trade is $c(X,\theta)$, where θ is a set of parameters to be estimated. Finally, let $p*$ be the price the dealer offers to the seller.

The dealer is risk-neutral with indirect utility function $E(p \mid X) - c(X,\theta) - p*$, the expected profit from purchasing the bond today for $p*$, then reselling it in the future for an expected price of $E(p \mid X)$, and incurring the expected intermediation cost, $c(X,\theta)$. The seller is also risk-neutral with indirect utility function $p* - v$, the price she receives for selling the bond today, $p*$, less her reservation value for the bond, v.

The seller and dealer engage in an alternating offer game with the possibility of breakdown, the solution of which can be described by the generalized Nash solution. Let ρ be the bargaining power of the dealer relative to that of the seller, where $\rho \in [0,1]$. If $\rho = 0$, the seller has all the bargaining power, and if $\rho = 1$, the dealer has all the bargaining power.

The equilibrium transaction price $p*$ maximizes the generalized Nash product

$$\max_{p*} [E(p \mid X) - c(X,\theta) - p*]^{\rho}(p* - v)^{1-\rho} \tag{1}$$

subject to the participation constraints

$$E(p \mid X) - c(X,\theta) - p* \geq 0, \tag{2}$$

$$p* - v \geq 0. \tag{3}$$

Condition (2) requires the price to be less than the dealer's expected net revenues from reselling the bond, and condition (3) requires the price to exceed the seller's reservation value. The participation constraints can only be satisfied if there are positive gains from trade:

$$E(p \mid X) - c(X,\theta) - v \geq 0. \tag{4}$$

If the gains from trade are not positive, the game ends and no trade takes place.

The first-order condition when the gains from trade are positive is

$$(1 - \rho)[E(p \mid X) - c(X,\theta) - p*] + \rho(v - p*) = 0. \tag{5}$$

Solving (5) for $p*$, the equilibrium offer price is

$$p* = \rho v + (1 - \rho)[E(p \mid X) - c(X,\theta)]. \tag{6}$$

The transaction price is a weighted average of the seller's reservation value and the dealer's expected net revenues from reselling the bond. The weights are given by the relative bargaining power of the counterparties.

If the seller has all the bargaining power ($\rho = 0$), the offer price is equal to the expected cash flows that the dealer receives from reselling the bond, $E(p \mid X) - c(X,\theta)$, and the dealer makes zero expected profits. If the dealer has all the bargaining power ($\rho = 1$), the offer price is equal to the seller's reservation price for the bond and the dealer makes positive expected profits.

Once the dealer has the bonds she will in turn sell them for p, which will differ from the expected price by an expectational error with mean zero:

$$e \equiv p - E(p \mid X). \tag{7}$$

The realized markup will consist of this expectational error, along with that portion of the surplus the dealer is able to extract from the seller. To see this, note that

$$\begin{aligned} p - p* &= E(p \mid X) + e - p* \\ &= E(p \mid X) + e - \{\rho v + (1-\rho)[E(p \mid X) - c(X,\theta)]\} \\ &= c(X,\theta) + e + \rho[E(p \mid X) - c(X,\theta) - v]. \end{aligned} \tag{8}$$

A transaction takes place, and $p - p*$ is observed, if and only if

$$\rho[E(p \mid X) - c(X,\theta) - v] \geq 0. \tag{9}$$

Condition (9) follows because $\rho \in [0,1]$ and because the participation constraints in Equations (2) and (3) can both be satisfied if and only if the gains to trade are positive: $E(p \mid X) - c(X,\theta) - v \geq 0$.

3.2 Estimating the model

The model in Equation (8) fits naturally into a specification that can be estimated using stochastic frontier analysis. Greene (2002) and Kumbhakar and Lovell (2003) provided descriptions of stochastic frontier analysis at a textbook level. The classic applications of stochastic frontier analysis in production, such as Aigner, Lovell, and Schmidt (1977), estimate production or cost functions that are viewed as the most efficient outcomes possible. Individual observations deviate from this ideal by a symmetric error that has zero mean and by a one-sided error that is interpreted as inefficiency specific to that firm. Stochastic frontier analysis has been applied in financial economics by Hunt-McCool, Koh, and Francis (1996) and Koop and Li (2001) to study IPO underpricing, by Berger and Mester (1997) and Altunbas et al. (2001) to study efficiency in the banking industry, and by Habib and Ljungqvist (2005) to study the role of incentives in mitigating agency costs.

In our application, the cost of intermediating the trade to the dealer, $c(X_i,\theta)$, can be viewed as the "efficient" markup from the standpoint of seller i, where we now use i to index the transactions in the data. Such a markup would be attained on average by sellers who have sufficient bargaining power to drive dealers to their reservation prices. Even in that case, Equation (8) implies that observed markups would vary around this cost by an expectational error, e_i. If sellers lack bargaining power relative to dealers, the observed markup will also deviate from the dealer's cost by a one-sided error that reflects the dealer's market or bargaining power and the seller's reservation price. Measures of the relative size, or variance, of the expectational and one-sided errors provide information about the relative importance of the dealers' market power.

It is natural to interpret the markup in the model in percentage terms. For trades taking place across days, we also interpret the model as applying to excess returns over what would be earned on an index of municipal bonds. Let $R_{index,i}$ be the return on a municipal bond index, and define the forecast error η_i

$$\eta_i = R_{index,i} - E(R_{index,i} \mid X_i). \tag{10}$$

To rewrite the model in percentage terms, divide Equation (8) by the initial bond price, p_i^*, and subtract the return on the municipal bond index:

$$\frac{p_i - p_i^*}{p_i^*} - R_{index,i} = \left[\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i} \mid X_i)\right] + \epsilon_i + \xi_i, \tag{11}$$

with

$$\epsilon_i \equiv \frac{e_i}{p_i^*} - \eta_i, \tag{12}$$

and

$$\xi_i \equiv \frac{\rho_i[E(p_i \mid X_i) - c(X_i,\theta) - v_i]}{p_i^*}. \tag{13}$$

Here $\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i}|X_i)$ are the dealer's costs in excess of the expected return on an index of municipal bonds. We will refer to $\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i}|X_i)$ as the cost of intermediation. Because p_i^* is positive, the normalization in Equation (13) does not change the one-sided nature of the second error term.

To arrive at an econometric specification, we put parametric structure on the distribution of the error, ϵ_i, and the distribution of the seller's reservation value, v_i. We assume $\epsilon_i \sim \mathcal{N}(0,\sigma_i^2)$. We allow the distribution to depend on the characteristics of the bond or the seller through the variance. Reservation values are drawn from a distribution centered on the dealer's expected net proceeds from selling the bond. That is,

$$E(v_i \mid X_i) = E(p_i \mid X_i) - c(X_i,\theta). \tag{14}$$

It is natural to assume that potential sellers' reservation values are rational in this sense. They reflect correct expectations on average. The assumption might be viewed as unrealistic, in that half of the potential arriving customers fail to sell their bonds. Given the institutional setting, however, the set of potential sellers, interpreted broadly, is unlikely to be observable. The model could be generalized to allow for some other truncation point in the distribution, but without data on arrivals, it is unclear how to identify an appropriate cutoff.

Define ν_i as the deviation of v_i around its mean normalized by the price:

$$\nu_i \equiv \frac{E(v_i \mid X_i) - v_i}{p_i^*}. \tag{15}$$

Substituting Equation (15) into Equation (13),

$$\xi_i = \rho_i \nu_i. \tag{16}$$

We parameterize ν_i as a double exponential, where the density on each side of the mean is 0.5 times the exponential density with parameter λ_i. Again, we allow the distribution from which the reservation price is drawn to depend on the characteristics of the trade. The distribution conditional on $\nu_i \geq 0$ is exponential with parameter λ_i. The one-sided error ξ_i, conditional on $\xi_i \geq 0$, therefore is exponentially distributed with parameter (λ_i/ρ_i). The first two moments of ξ_i are

$$E(\xi_i \mid \xi_i \geq 0) = \frac{\rho_i}{\lambda_i}, \tag{17}$$

$$\mathrm{Var}(\xi_i \mid \xi_i \geq 0) = \left(\frac{\rho_i}{\lambda_i}\right)^2. \tag{18}$$

A natural specification is to allow the distribution to depend on features such as the par value of the trade and how the trade is handled by the dealer. We would expect large trades to originate with institutional investors, who are likely to be better informed about market conditions and who, because of the repeat business they bring to dealers, can bargain more effectively over any given transaction. We might also expect dealers to accept trades anticipating the involvement of other dealers, or planning to split the bond between multiple buyers, only in situations where they anticipate a large share of significant gains to trade.

We estimate specifications where the exponential error term has parameter (ρ_i/λ_i) that is a log-linear function of the explanatory variables $Z_i = \{Z_{ik}\}_{k=1}^{K}$:

$$\frac{\rho_i}{\lambda_i} = a_0 \prod_{k=1}^{K} e^{a_k Z_{ik}}. \tag{19}$$

The explanatory variables include the natural logarithm of trade size (par value), and dummies for whether the trade is split, or moves through other dealers. We also allow the standard deviation of the symmetric error, σ_i, to depend on the same variables:

$$\sigma_i = b_0 \prod_{k=1}^{K} e^{b_k Z_{ik}}. \tag{20}$$

If $a_0 = 0$, then $(\rho_i/\lambda_i) = 0$; the sellers have all the bargaining power, and the realized markup does not contain an asymmetric error. If $a_0 \neq 0$, then $(\rho_i/\lambda_i) \neq 0$; the dealers have bargaining power, and the realized markup contains an asymmetric error. Variables that increase (ρ_i/λ_i) suggest higher bargaining power for the dealers.

Finally, we parameterize the cost of intermediation as a linear function of the explanatory variables:

$$\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i} \mid X_i) = \theta_0 + \sum_{l=1}^{L} \theta_l X_{il} \tag{21}$$

where L is the number of regressors. Our methods also allow for the parties to condition their decisions on unobserved heterogeneity in the cost function.[9]

We report the estimates of the stochastic frontier models for both narrowly and more broadly defined samples. The general conclusions of our analysis are very robust to the selection criteria we employ.

3.3 Issues of interpretation and identification

The parameters governing the dealer's bargaining power, ρ_i, and the size of the asymmetric error, λ_i, are not separately identified in the model. This corresponds to the underlying economics. We would expect more sophisticated customers to come to the market with a better understanding of the dealers' costs and resale opportunities and to therefore demand

[9] The model can be generalized to allow the cost function to consist of two parts: $c^*(X_i,\theta) + w_i$, where the conditioning variables in X_i are observed by the econometrician, and w_i is a mean-zero source of heterogeneity observed by both the buyer and the seller, but not the econometrician. Then replacing $c(X_i,\theta)$ in Equations (6)–(11) with $c^*(X_i,\theta) + w_i$ gives

$$\frac{p_i - p_i^*}{p_i^*} - R_{index,i} = \left[\frac{c^*(X_i,\theta)}{p_i^*} + \frac{w_i}{p_i^*} - E(R_{index,i} \mid X_i)\right] + \epsilon_i + \xi_i$$

where, now,

$$\xi_i \equiv \frac{\rho_i[E(p_i \mid X_i) - c^*(X_i,\theta) - w_i - v_i]}{p_i^*}.$$

We can now interpret $\frac{w_i}{p_i^*} + \epsilon_i$ as the symmetric portion of the error in the model, and assuming that

$$E(v_i \mid X_i) = E(p_i \mid X_i) - c^*(X_i,\theta) - w_i,$$

we can interpret ξ_i as the asymmetric portion of the error and parameterize each as in the body of the article.

prices leading to lower dealer's profits. This could be interpreted either as a more concentrated distribution for their reservation prices, or as greater bargaining power for the customer and thus heterogeneity in ρ_i.

The model ascribes the dealer's bargaining power to his dealings with the customers seeking liquidity by selling their bonds. Whatever rents the dealer can extract from the eventual buyer are subsumed in the expectations of the future sales price. This may seem natural in this institutional setting, where the seller has a unique object to dispose of, while a buyer can choose among many close substitutes. We cannot, however, interpret the asymmetric error as measuring a cost borne entirely by the seller.

Our methods estimate the model using pooled cross-sectional and time-series data. Because the errors in the model may reflect common time-series effects not captured by the adjustment for index returns, and common cross-sectional correlation for similar types of bonds, our standard errors will likely be understated. Practically speaking, however, we have in the full sample over 4.5 million observations, so these effects seem unlikely to be of great importance. The results are robust to stratifying our sample into rising, falling, and steady markets, to the sample selection criteria, and to including other explanatory variables.

The stochastic frontier model is identified through functional form. From Figure 1 it is obvious that the distribution of markups is highly skewed and that both this skewness and the variance change together with the size of the trade. The figure plots the unconditional distribution. In principle, it is possible that this behavior disappears once we condition on the variables entering the cost function, in which case there will be little for the one-sided error to explain. Assuming that is not the case, the model ascribes to dealer's market power variation that can be captured by a term with a mean that is equal to its standard deviation. Remaining variation is attributed to the cost function and symmetric error. This is precisely the structure predicted by the model. As in any structural setting, however, there may be alternative explanations for the conditional skewness the econometric specification captures. Such an alternative explanation is not apparent to the authors, but it would presumably involve some type of variation in costs or prices that is highly skewed and is unobservable at the time of the initial purchase or omitted from the specification.

Of course, more than just the first two moments are involved in identification of the model. Define the empirical residuals for the model

$$\hat{f}_i = \frac{p_i - p_i^*}{p_i^*} - R_{index,i} - \hat{\theta}_0 + \sum_{l=1}^{L} \hat{\theta}_l X_{il}. \tag{22}$$

If the residual had no one-sided component, then θ would be estimated using the standard OLS moment condition:

$$E[f_i X_i'] = 0. \tag{23}$$

In Appendix C, we show that instead, the maximum-likelihood estimator uses the moment conditions

$$E[f_i] = \frac{\rho}{\lambda}, \tag{24}$$

$$E\left[\frac{\phi\left(\frac{f_i}{\sigma} - \sigma\left(\frac{\lambda}{\rho}\right)\right)\sigma f_i}{\Phi\left(\frac{f_i}{\sigma} - \sigma\left(\frac{\lambda}{\rho}\right)\right)}\right] = E[E(\epsilon_i|f_i)] = 0, \tag{25}$$

and

$$E\left[\left(-\frac{\phi\left(\frac{f_i}{\sigma} - \sigma\left(\frac{\lambda}{\rho}\right)\right)\frac{1}{\sigma}}{\Phi\left(\frac{f_i}{\sigma} - \sigma\left(\frac{\lambda}{\rho}\right)\right)} + \left(\frac{\lambda}{\rho}\right)\right)X_i'\right] = 0, \tag{26}$$

where ϕ is the standard normal density and Φ is the standard normal distribution function. Condition (24) is that the unconditional mean of the residuals equals the mean of the exponential, (ρ/λ) and condition (25) is that the unconditional mean of the normal part of the errors equals zero. The term in large parentheses in (26) is a generalized residual for the model. This nonlinear function of the model's residuals should be orthogonal to the explanatory variables—the condition holds if the density of the residuals is that of the sum of independent normal random variable and an exponential random variable.

4. Empirical Results

Table 9 reports the results of fitting the empirical model in Equations (19)–(21), conditioning on the variables defined in Table 2. All the variables have been constructed to include only information available at the time of the initial purchase from a customer. The model is estimated using the full sample. It includes the immediate matches, the round-trip transactions, and all transactions that can be identified with the FIFO method. The specification allows the variances of both the symmetric and the asymmetric errors to depend on the size of the transaction. We use daily returns on the Lehman Brothers Municipal Bond Index as the index return, R_{index}. Because we use daily returns, $R_{index} = 0$ if both sides of the transaction are completed on the same day.

To aid the interpretation of the coefficients, we transformed the independent variables by subtracting their sample means. The intercept term in the cost of intermediation function therefore measures the

Table 9
Parameter estimates of the stochastic frontier model: full sample

	No market power		Market power	
	Cost function parameters			
Constant	177.1	(0.1)	69.5	(0.1)
$\ln(Par) \times Par \in (0{,}100k)$	−44.4	(0.1)	−32.8	(0.1)
$\ln(Par) \times Par \in [100k{,}500k)$	−39.2	(0.1)	−22.8	(0.2)
$\ln(Par) \times Par \in [500k{,}\infty)$	−16.5	(0.2)	−5.3	(0.2)
Maturity	3.7	(0.0)	2.9	(0.0)
Coupon	7.9	(0.1)	5.5	(0.1)
Age	0.0	(0.0)	−0.7	(0.0)
Issue size	−9.8	(0.1)	−11.0	(0.1)
Transaction frequency	−36.3	(0.3)	−33.1	(0.2)
State volume	0.0	(0.3)	1.2	(0.3)
State order imbalance	3.5	(2.1)	3.0	(1.9)
State new issues	29.2	(0.6)	25.8	(0.6)
Treasury yield	7.9	(0.1)	6.0	(0.1)
Treasury volatility	61.5	(1.2)	57.2	(1.1)
Coupon zero or missing	37.6	(0.5)	6.0	(0.5)
Rating available	32.0	(0.2)	29.8	(0.2)
Rating A-BBB	63.7	(0.4)	45.7	(0.4)
Rating BB-B	126.9	(1.2)	72.6	(1.2)
Rating CCC-D	132.2	(1.7)	47.3	(1.8)
Callable bond	19.2	(0.2)	22.6	(0.2)
Taxable bond	−43.4	(0.5)	−57.4	(0.5)
Pre-refunded bond	−52.5	(0.6)	−47.5	(0.5)
	Symmetric error distribution			
Constant	168.4	(0.0)	124.7	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in (0{,}100k)$	−3.5	(0.0)	−1.1	(0.0)
$\ln(Par) \times 10^{-2} \times Par \in [100k{,}500k)$	−4.7	(0.1)	−0.7	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in [500k{,}\infty)$	−3.9	(0.1)	−2.3	(0.1)
	Asymmetric error distribution			
Constant	–	–	106.2	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in (0{,}100k)$	–	–	−9.2	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in [100k{,}500k)$	–	–	−12.6	(0.2)
$\ln(Par) \times 10^{-2} \times Par \in [500k{,}\infty)$	–	–	−9.8	(0.3)
Difference in the log-likelihood		184,807		
Observations		4,583,588		

This table reports maximum-likelihood estimates from fitting the model

$$\frac{p_i - p_i^*}{p_i^*} - R_{index,i} = \theta_0 + \sum_{l=1}^{L} \theta_l X_{il} + \epsilon_i + \xi_i,$$

with p_i the price that the dealer sells the bond for, p_i^* the price that the dealer pays the seller for the bond, X_{il} for $l = 1,\ldots,L$ conditioning variables. The residual ϵ_i is normally distributed with standard deviation $b_0 \prod_{k=1}^{K} e^{b_k Z_{ik}}$, and ξ_i is exponential distributed with parameters $a_0 \prod_{k=1}^{K} e^{a_k Z_{ik}}$, with Z_{ik} for $k = 1,\ldots,K$ conditioning variables. The cost function parameters θ_l, $l = 0,\ldots,L$, the coefficients b_k, $k = 0,\ldots K$, in the symmetric error distribution, and the coefficients a_k in the asymmetric error distribution are reported. The dichotomous dummy explanatory variables listed in Table 2 are included in the cost function specification, but the coefficient estimates are omitted from the table. The dependent variable is measured in basis points, and the explanatory variables are demeaned. The first two columns report the parameter estimates and standard errors in parentheses for the model with $a_0 = a_1 = \ldots = a_K = 0$, and the second two columns report the parameter estimates and standard errors in parentheses for the unrestricted model.

expected cost of a typical trade in basis points. The intercept term in the volatility of the symmetric error measures the volatility of the symmetric error for a typical trade, and the intercept term in the exponential error term measures the expected profit from market power, measured in basis points.

The first two columns of Table 9 report the results of fitting the model with all the parameters in the one-sided error constrained to zero.[10] With the one-sided error constrained to zero, the sellers have all the bargaining power, and the model is a standard, pooled cross-sectional, and time-series regression analysis of the determinants of the costs of trading. The model is estimated by maximum-likelihood. Table 9 reports the coefficients on the cost of intermediation function, $\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i} \mid X_i)$ and the coefficients in the residual standard deviation.

Almost all the explanatory variables are estimated with high precision, reflecting the sample size of 4.5 million observations. The signs are generally consistent with higher markups where one would expect to see them if dealers do demand compensation for bearing liquidity costs and interest rate risk. Within each size category, trade size lowers the gains to the dealers. In environments with high interest rates, or volatile interest rates, markups are larger. Bonds with higher liquidity, measured by transaction frequency and issue size, earn dealers lower markups. Pre-refunded bonds, which are free of default risk because they are backed by trusts funded with treasury bonds, earn much lower spreads. Lower rated bonds trade with much bigger markups.

The second two columns of Table 9 report maximum-likelihood estimates of a stochastic frontier model with specifications (19)–(21), with the parameters unrestricted. Table 9 reports the coefficients in the linear model for the "efficient" markup, or, $\frac{c(X_i,\theta)}{p_i^*} - E(R_{index,i} \mid X_i)$, the parameter estimates for the two-sided and one-sided error distributions, and the difference in the log-likelihoods from including the one-sided error.

Because the explanatory variables are demeaned in the model, the intercept terms are estimates of the expected efficient cost and the expected loss to dealer's market power on an average transaction. The constant in the estimates for the asymmetric error (106.2 basis points) is much larger than the constant in the cost function (69.5 basis points). The constant in the cost function has a natural interpretation of the average efficient cost of search or intermediation. Much of the markup on the average transaction is attributed by the model to the dealer's advantage in bargaining power. Similarly, the intercept in the cost function falls by more than half, and the constant in the volatility of the symmetric error falls by a quarter when the model allows for the

[10] To save space, we do not report all the coefficients and standard errors on the dummy variables.

one-sided error. More than half of the expected costs of trading and much of the conditional volatility of the trading costs are ascribed to the dealers' market power.

Several of the coefficients in the one-sided error specification decrease in absolute value relative to the model with only a symmetric error distribution. Much of the effect of size, and also of bond rating, is subsumed by allowing for dealer's bargaining power. For example, bonds with the lowest rating appear to cost an extra 1.3% to handle when the model does not allow for dealer's market power. This falls to less than half a percent in the unconstrained model. That the effects of these variables are in large part subsumed by the asymmetric error term suggests that their explanatory power for the markups may be due to the ability of dealers to extract greater rents in these types of trades rather than their effect on the dealer's costs. The mean and variance of the asymmetric error drop with the size of the transaction. The model ascribes more bargaining power to dealers in their interactions with smaller traders.

Table 10 summarizes that the broad, qualitative features of the results are robust to how the sample is selected. The table reports constrained and unconstrained parameter estimates for the immediate matches, the additional transactions selected by the round-trip procedure, and the additional observations identified by the FIFO method. Together, the three samples make up our full sample.

In every case the intercept term in the cost function falls dramatically under the stochastic frontier specification, as does the impact of size. Evidently, the "efficient cost frontier" identified by the model is much lower than the costs investors actually face. The size of the symmetric error falls in every case in the stochastic frontier specification. Its size relative to the asymmetric error, however, differs. The symmetric error is much smaller than the asymmetric error for the immediate matches, the symmetric error about the same size for the round-trip transactions, and is larger for the FIFO sample. This is what we would expect given the differences in the time the bonds are in inventory for the three cases (Table 5).[11]

[11] Using the full sample, we also estimated the model separately for days with rising markets, falling markets, and mixed movements in the yield curve. Rising markets are defined as days t on which four points on the yield curve (3-month, 5-, 10-, and 30-year terms) decrease between the close at $t-1$ and $t+1$. Mixed markets are days on which the yield curve does not parallel shift. Falling markets are associated with upward shifts of the yield curve at all maturities. This stratification has virtually no effect on the results. The intercept terms in the cost functions, for example, vary between 62 and 77 basis points. The constants in the asymmetric error terms vary between 104 and 104, and in the symmetric error term between 122 and 126. The coefficients on the size terms are extremely similar across the three sets of estimates. We have also experimented with using a half-normal distribution for the one-sided error. Our empirical findings are robust to such a change. Like the estimates in Tables 9 and 10 the intercept term in the cost function is much lower than that in the asymmetric error specification. The likelihood ratios associated with allowing for the asymmetric error are also very large regardless of the specification used.

Table 10
Parameter estimates of the stochastic frontier model: full sample, stratified by sample selection

	Immediate matches				Round-trip transactions				FIFO transactions			
	No market power		Market power		No market power		Market power		No market power		Market power	
	Cost function parameters											
Constant	140.4	(0.1)	60.7	(0.2)	204.1	(0.1)	113.2	(0.2)	174.7	(0.2)	36.2	(0.3)
$\ln(Par) \times Par \in (0{,}100k)$	−41.8	(0.1)	−24.8	(0.1)	−43.8	(0.1)	−32.8	(0.1)	−40.9	(0.1)	−25.4	(0.2)
$\ln(Par) \times Par \in [100k{,}500k)$	−30.1	(0.2)	−3.8	(0.1)	−44.9	(0.2)	−29.7	(0.2)	−28.1	(0.2)	−7.1	(0.3)
$\ln(Par) \times Par \in [500k{,}\infty)$	−5.6	(0.1)	5.2	(0.1)	−47.7	(0.3)	−27.1	(0.4)	−5.2	(0.2)	5.0	(0.3)
	Symmetric error distribution											
Constant	89.2	(0.1)	43.7	(0.1)	125.8	(0.1)	85.2	(0.1)	217.3	(0.1)	164.5	(0.2)
$\ln(Par) \times 10^{-2} \times Par \in (0{,}100k)$	−20.3	(0.1)	−28.4	(0.1)	−7.9	(0.0)	−5.7	(0.1)	−7.8	(0.0)	−7.0	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in [100k{,}500k)$	−26.9	(0.2)	−6.8	(0.4)	−8.7	(0.1)	−6.0	(0.2)	−7.8	(0.1)	−3.1	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in [500k{,}\infty)$	−29.1	(0.2)	−33.1	(0.3)	3.6	(0.2)	9.1	(0.2)	−5.2	(0.1)	−3.3	(0.1)
	Asymmetric error distribution											
Constant	–	–	69.9	(0.2)	–	–	84.5	(0.2)	–	–	138.5	(0.3)
$\ln(Par) \times 10^{-2} \times Par \in (0{,}100k)$	–	–	−24.3	(0.2)	–	–	−13.4	(0.1)	–	–	−11.1	(0.1)
$\ln(Par) \times 10^{-2} \times Par \in [100k{,}500k)$	–	–	−39.2	(0.3)	–	–	−17.6	(0.3)	–	–	−15.7	(0.3)
$\ln(Par) \times 10^{-2} \times Par \in [500k{,}\infty)$	–	–	−69.5	(0.5)	–	–	−29.8	(1.1)	–	–	−10.1	(0.4)
Difference in log-likelihood	65,697				121,648				52,617			
Observations	747,636				1,961,217				1,874,735			

This table reports maximum-likelihood estimates separately for the three subsamples corresponding to immediate matches, round-trip transactions, and the FIFO sample. The cost function parameters θ_l, $l = 0,\ldots,L$, the coefficients b_k, $k = 0,\ldots,K$, in the symmetric error distribution, and the coefficients a_k in the asymmetric error distribution are reported. Additional explanatory variables are included in the cost function, but the coefficient estimates are omitted from the table. These explanatory variables are rating available, rating **A-BBB**, rating **BB-B**, rating CCC-D, maturity, coupon, coupon zero or missing, age, transaction frequency, issue size, state volume, state order imbalance, state new issues, treasury yield, treasury volatility, and as in Table 9 all dichotomous variables specifying the type of bond described in Table 2. The dependent variable is measured in basis points, and the explanatory variables are demeaned. The first two columns report the parameter estimates and standard errors in parentheses for the model with $a_0 = a_1 = \ldots = a_K = 0$, and the second two columns report the parameter estimates and standard errors in parentheses for the unrestricted model.

5. Trading Costs and the Form of Intermediation

A dealer makes many choices about how to sell bonds subsequent to, or simultaneously with, purchasing them. From our sample, we know if the dealer splits the order into smaller blocks or not, and if the order goes through other dealers or directly to customers. The results in Tables 6–8 show that the realized markups depend on how the dealers process the bond. To the extent that these decisions are predictable at the time of the initial purchase, the cost function and the measure of dealer's market power might depend on how the parties expect that the trade will be intermediated.

Several questions of interest about dealer choices can be addressed by incorporating these decisions into our model. Are the decisions dealers make about the form of intermediation based on information at the time of the initial purchase, and are they entirely an ex-post response? For example, dealers may decide to split orders or involve other dealers only when, through new information, they find they can reach high valuation buyers in these ways, or when market conditions turn against them unexpectedly, and they are forced to seek liquidity from additional counterparties. If the decisions are an ex post response of this sort, they should be difficult to predict, and predictions of them should have little impact on our measure of expected costs or dealer's bargaining power.

Alternatively, there may be certain types of trades that have high anticipated costs to dealers, because they expect to employ more layers of intermediation. In this case, the predicted decisions should have explanatory power in our estimated "efficient" cost function. It is also possible that dealers choose to undertake extra layers of intermediation in situations where they anticipate that it will allow them to reach high valuation investors at low cost. Splitting of blocks or interdealer trades might occur largely for complex, illiquid bonds. Dealers may have more bargaining power in these trades, because more sophistication and knowledge of market infrastructure is required. Under these circumstances, a trade that is more likely to involve complex intermediation should increase the size of the asymmetric error in our stochastic frontier model.

To address these questions, we estimate the stochastic frontier model conditioning the cost function and the distributions of the error terms on estimates of the probability the trade is intermediated in different ways. For each transaction define the discrete random variable d_i:

$$d_i = \begin{cases} 1, & \text{if the dealer does not split and does not sell to another dealer,} \\ 2, & \text{if the dealer splits the order and does not sell to another dealer,} \\ 3, & \text{if the dealer does not split the order and sells to another dealer,} \\ 4, & \text{if the dealer splits the order and sells to another dealer.} \end{cases} \tag{27}$$

Let D_{ij} be the indicator for $d_i = j$. We use a multinomial logit model to estimate

$$\Pr(d_i = j|W_i) = E[D_{ij}|W_i], \tag{28}$$

where W_i includes only information available at the time of the initial purchase from a customer. We then allow the expected costs, the standard deviation of the symmetric error, and the parameters in the distribution of the one-sided error to all depend on the probabilities. Because the conditional probabilities are generated regressors, we construct unbiased standard errors using the nonparametric bootstrap procedure. We resample one tenth of the observations with replacement from the data and re-estimate both the discrete choice equation and the markup equation, respectively to form bootstrap standard errors.

We estimated separate multinomial logit models for each of three trade-size categories. The results (which are not tabulated for reasons of space) indicate that the decisions to split trades and involve other dealers depend, as one might expect, on the bond's characteristics, such as rating, on volume and liquidity measures, and on transaction size. There is also evidence of serial dependence in these decisions. It is more likely that a trade will be split or moved through other dealers when recent trades were similarly handled. The estimates from the logit models suggest that the decisions are predictable; for all par values, we reject the null hypotheses that the conditioning variables, W_i, do not predict the decisions.[12]

The models in Table 11 include the conditional probabilities of the intermediation services interacted with the transaction size in both the cost function and the specifications of the error terms. Interacting the size dummy variables with the other variables is both empirically and economically important, because trade size proxies for heterogeneity in customer type. For instance, one would expect that pre-negotiation is more pervasive for larger trades. To interpret the coefficients in the specification for the asymmetric error term, note that they measure the effect of a change in the probability of different types of intermediation on the expected markup in percent, because they estimate $\frac{\partial(\frac{\rho}{\lambda})}{\partial E[D_{ij}|W_i]} \times \frac{1}{(\frac{\rho}{\lambda})}$. For example, for a large trade with par value over \$500,000, a 1% increase in the probability of a split with no interdealer trading is associated with a 1.74% increase in the expected gain to the dealer.

Our general results are robust to including information on how orders are processed. Allowing for an asymmetric error reduces the intercept in the cost function by over 50% and also decreases the estimated standard

[12] For $Par \in (0,100k)$: $\chi^2_{138} = 1{,}079{,}373$, p-value $= 0.00$, for $Par \in [100k,500k)$: $\chi^2_{138} = 314{,}684$, p-value $= 0.00$, and for $Par \in [500k,\infty)$: $\chi^2_{138} = 421{,}166$, p-value $= 0.00$.

Table 11
Parameter estimates of the stochastic frontier model: expected dealer choices

	No market power		Market power	
	Cost function parameters			
Constant	177.1	(0.3)	80.9	(1.1)
$\widehat{E}$(Split × no Interdealer) × *Par* ∈ (0,100*k*)	211.5	(6.6)	145.0	(11.0)
$\widehat{E}$(Split × no Interdealer) × *Par* ∈ [100,500*k*)	172.5	(7.4)	−5.3	(9.0)
$\widehat{E}$(Split × no Interdealer) × *Par* ∈ [500*k*,∞)	114.9	(7.5)	−25.3	(11.2)
$\widehat{E}$(No Split × Interdealer) × *Par* ∈ (0,100*k*)	331.7	(9.0)	239.7	(14.1)
$\widehat{E}$(No Split × Interdealer) × *Par* ∈ [100*k*,500*k*)	211.4	(11.7)	31.0	(20.4)
$\widehat{E}$(No Split × Interdealer) × *Par* ∈ [500*k*,∞)	253.2	(5.6)	106.2	(10.9)
$\widehat{E}$(Split × Interdealer) × *Par* ∈ (0,100*k*)	176.0	(8.8)	11.9	(10.9)
$\widehat{E}$(Split × Interdealer) × *Par* ∈ [100*k*,500*k*)	463.9	(78.6)	173.8	(57.9)
$\widehat{E}$(Split × Interdealer) × *Par* ∈ [500*k*,∞)	273.4	(10.5)	91.8	(14.2)
	Symmetric error distribution			
Constant	161.6	(0.3)	125.8	(0.5)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ (0,100*k*)	−6.5	(3.6)	−13.7	(5.8)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	92.2	(2.8)	67.8	(4.3)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	88.8	(3.4)	75.2	(5.1)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ (0,100*k*)	−4.0	(6.8)	−8.9	(7.1)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	−25.8	(7.0)	−98.4	(14.0)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	47.0	(2.7)	21.9	(5.2)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ (0,100*k*)	108.0	(4.2)	88.6	(5.5)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	144.9	(28.2)	82.8	(28.5)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	37.6	(4.4)	0.3	(5.3)
	Asymmetric error distribution			
Constant	–	–	88.8	(2.2)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ (0,100*k*)	–	–	108.7	(23.0)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	–	–	153.2	(7.9)
$\widehat{E}$(Split × No Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	–	–	174.5	(12.1)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ (0,100*k*)	–	–	111.5	(22.2)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	–	–	117.6	(24.1)
$\widehat{E}$(No Split × Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	–	–	156.3	(17.4)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ (0,100*k*)	–	–	270.7	(22.8)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ [100*k*,500*k*)	–	–	263.5	(28.9)
$\widehat{E}$(Split × Interdealer) ×10^2× *Par* ∈ [500*k*,∞)	–	–	221.4	(26.1)
Difference in log-likelihood function		148,258		
Observations		4,583,588		

This table reports maximum-likelihood estimates from fitting the model

$$\frac{p_i - p_i^*}{p_i^*} - R_{index,i} = \theta_0 + \sum_{j=2}^{J} \theta_{j-1} \hat{E}(D_{ij}|W_i) + \sum_{l=i}^{L} \theta_{J-1+l} X_{il} + \epsilon_i + \xi_i,$$

with p_i the price that the dealer sells the bond for, p_i^* the price that the dealer pays the seller for the bond, X_{il} for $l = 1,\ldots,L$ conditioning variables. The $j = 1,\ldots,J$ dichotomous variables D_{ij} represent the dealer's intermediation choices, and W_i is a set of conditioning variables known at the time of the initial transaction. The residual ϵ_i is normally distributed with standard deviation $b_0 \prod_{j=2}^{J} e^{b_{j-1}\widehat{E(D_{ij}|W_i)}} \prod_{k=1}^{K} e^{b_{J-1+k} Z_{ik}}$, and ξ_i is exponential distributed with parameters $a_0 \prod_{j=2}^{J} e^{a_{j-1}\widehat{E(D_{ij}|W_i)}} \quad \prod_{k=1}^{K} e^{a_{J-1+k} Z_{ik}}$, with Z_{ik} for $k = 1,\ldots,K$ conditioning variables.

The cost function parameters θ_j, $j = 0,\ldots,J-1$, the coefficients b_j, $j = 0,\ldots,J-1$, in the symmetric error distribution, and the coefficients a_j in the asymmetric error distribution are reported. Additional explanatory variables are included, but the coefficient estimates are omitted from the table. These explanatory variables are the same as in Table 9. The dependent variable is measured in basis points, and the explanatory variables are demeaned.

The first two columns report the parameter estimates and robust standard errors in parentheses for the model with $a_0 = a_1 = \ldots = a_{J+K-1} = 0$, and the second two columns report the parameter estimates and robust standard errors in parentheses for the unrestricted model. Robust standard errors are obtained from a nonparametric bootstrap. We use a total of 50 replications with replacement of one-tenth of the observations. In each replication we re-estimate both the discrete choice equation and the markup equation, respectively, and regenerate $\widehat{E}(D_{ij}|W_i) = \widehat{Pr}(d_i = j|W_i)$ using the corresponding first-stage estimates.

deviation of the symmetric error. The constant term in the asymmetric error term is substantial.

The coefficients on the conditional probabilities in the cost function measure the impact of greater probability of complex intermediation on the "efficient" round-trip cost of trading. Small trades appear to be particularly costly with a larger expectation that the trade will be split or moved through another dealer. The largest effect on the anticipated cost of trading for mid-sized customer sales to dealers is from the variable that predicts that a trade is both split and moved through multiple dealers. Overall, these results show that, as one would expect, more complex types of intermediation are costly. The one exception to this is lower cost when a very large trade is expected to be split. Presumably, very large trades that are amenable to splitting can be more efficiently disposed of in this manner.

All the conditional probabilities for higher levels of intermediation services increase the size of the asymmetric error, which, in the theoretical model, measures dealer's bargaining power. It is not surprising that these variables enter both the cost function and the asymmetric error terms with the same sign. Involving multiple counterparties is presumably more costly in time and effort, because more search, coordination, and processing are required. The types of trades that require these additional efforts may also be those where dealers have more leverage with customers. The access they enjoy to networks of other dealers and potential customers, and the expertise they possess, are scarce resources. They command higher rents in situations where those scarce resources are more essential.

A partial explanation for the result that involving other dealers increases the dealers' bargaining power may involve the types of dealers who tend to sell bonds to other dealers. Blocks of bonds purchased from institutions are likely to be split into smaller blocks to sell them to retail investors. Some large dealers have no retail distribution capacity. They will step forward to provide liquidity to an institution wishing to sell a large position knowing that they will probably need to sell the bonds to another dealer and appear to do so only when the terms are particularly advantageous. The gains to dealers may also reflect interdealer trading as a way of circumventing caps on dealer's profits in any one round-trip transaction.

6. Conclusion

Corporate and municipal bonds have traditionally been traded primarily in opaque, decentralized markets in which price information is costly to gather. Such a market structure may offer opportunities for intermediaries, in their role of facilitating the matching of buyers and sellers and providing liquidity, to extract monopoly rents or cross-subsidize across different groups of customers. But the extent to which they operate competitively is an empirical issue. Data on trade in these markets, however, have been difficult to obtain in the past, and even when such data are available separately, identifying costs and rents to the intermediary is problematic.

We analyze a comprehensive database of all trades between broker-dealers in municipal bonds and their customers. The data were only released to the public with a substantial lag, and thus the market remained relatively opaque to the traders themselves even though the data are available to researchers. Because the bonds trade infrequently, these transaction data are low frequency for each bond, but because there are so many bonds we have millions of observations.

We show that the pattern of dealer's profits in round-trip transactions implies that dealers earn lower average profits on larger trades, even while the dealers bear more risk on larger trades. We formulate a simple structural model that allows us to estimate measures of dealer's bargaining power and relate it to characteristics of the trades. The model is estimated as a stochastic frontier analysis model. The results suggest that dealers exercise substantial market power in their trades with customers. Our measures of market power decrease in trade size and appear larger for trades that are more complex to intermediate.

Appendix A: Data Filter

A number of transactions in the MSRB database are missing information about the bond issue. In case of missing bond features such as coupon rate, issue date, or maturity date, we perform searches over the entire database to recover the missing data. We use the bond description text field to identify the state in which the bond was issued, and the type and purpose of the bond. We identify the issuer from the CUSIP number. If the time stamp of the transaction is missing (or zero), we assume the transaction takes place at 6 a.m.

Next we eliminate obvious data errors such as negative prices, and maturity or issuance dates that do not make sense. We drop observations where the transaction size is missing, or where the reported par amount is less than 1,000. The remaining observations on the independent variables used in the estimation are winsorized at the 0.5 and 99.5% levels if they are symmetric, or at the 99.5% level if their values are one-sided.

MSRB reporting rules do not explicitly require broker-dealers to report both the transaction price and yield. In many instances no yield is reported, and in fewer cases the transaction price is missing. We drop observations where the transaction price is missing and truncate the distribution of net markups at 0.5 and 99.5%.

Appendix B: Sample Selection

The MSRB transaction database does not identify the broker-dealer who is intermediating a given transaction, nor the identity of the retail/institutional buyer or seller. The database, however, allows identifying the direction of the transaction—whether the transaction is a purchase from customer, a sale to customer, or an interdealer transaction—and the time of the transaction up to the minute.

We use a matching algorithm to select purchase and sales tickets that are likely to be associated with the same dealer. The basic principle is to match a given dealer purchase with consecutive sales to customers up until the next purchase from a customer or the end of the sample period. If the par amount of a sales ticket in this sequence is larger than the initial purchase amount, this particular sales transaction is ignored. All remaining sales transactions are aggregated. If the cumulative par amount of these sales transactions up until the next purchase equals the par amount of the preceding purchase, we have identified a matched pair. If we find a one-to-one match between a purchase and a single sale, we call these pairs "no splits." Otherwise the purchase is being split into several smaller units. We record the time of the purchase and the par weighted-average transaction time of the sales. In addition we determine whether transactions occurred between dealers in the same issue at any time between the purchase and sale date and mark the matched pair as involving interdealer intermediation.

We match the remaining purchase and sales transactions by applying a FIFO rule. After eliminating all transactions associated with matched pairs, we start with the first purchase in the remaining transaction sequence for each bond issue. We assign to it all consecutive sales transactions until they aggregate to or exceed the purchase amount—independently of whether other purchases occur in the meantime. If the par amount of these sales tickets exceeds that of the purchase, we allocate the residual to the next purchase if it occurred before the last of the consecutive sales. Otherwise we discard the residual. We record a partial match, if the par amount of the purchase exceeds that of all sales until the end of the sample. The matching process continues with the next purchase until the purchase is filled.

Using these matching rules we construct three separate samples. Tables 4 and 5 describe the different sampling criteria. We calculate net markups as the difference between the realized gross markup in percent and the index return on a maturity-matched portfolio of municipal bonds. The gross markup is the relative difference between the dealer's purchase price and the par weighted-average sale price:

$$\text{Net markup} = \text{Gross markup-Return on index}, \tag{B1}$$

$$\text{Gross markup} = \frac{\text{Par-weighted average sale price-purchase price}}{\text{Purchase price}}. \tag{B2}$$

Appendix C: Identification

This appendix discusses identification issues for our empirical model. The model is

$$y_i = X_i\theta + \epsilon_i + \xi_i, \tag{C1}$$

with ϵ_i normally distributed with standard deviation σ and ξ_i exponentially distributed with parameter λ/ρ, both random variables independent of each other:

$$\epsilon_i \sim \mathcal{N}(0,\sigma^2), \quad \xi_i \sim \left(\frac{\lambda}{\rho}\right) e^{-\left(\frac{\lambda}{\rho}\right)\xi}. \tag{C2}$$

Define $f_i = \epsilon_i + \xi_i$, we need the density of f_i for the log-likelihood. Using the standard formula for the cumulative distribution function (cdf) for sums of independent random variables,

$$\Pr(f_i \leq f) = \int_0^\infty \Phi\left(\frac{f-s}{\sigma}\right)\left(\frac{\lambda}{\rho}\right) e^{-s\left(\frac{\lambda}{\rho}\right)} ds, \tag{C3}$$

with Φ the standard normal cdf. Letting ϕ be the standard normal density, the density function for f_i is

$$\frac{1}{\sigma}\int_0^\infty \phi\left(\frac{f-s}{\sigma}\right)(\lambda/\rho)e^{-s(\lambda/\rho)}ds. = \int_0^\infty \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(f-s)^2}{2\sigma^2}}(\lambda/\rho)e^{-(\lambda/\rho)s}ds$$
$$= \int_0^\infty \frac{1}{\sqrt{2\pi}\sigma}(\lambda/\rho)e^{-\frac{f^2+s^2-2fs+2\sigma^2(\lambda/\rho)}{2\sigma^2}}ds$$
$$= \int_0^\infty \frac{1}{\sqrt{2\pi}\sigma}(\lambda/\rho)e^{-\frac{[s-[f-\sigma^2(\lambda/\rho)]]^2}{2\sigma^2}-(\lambda/\rho)f+\frac{1}{2}\sigma^2(\lambda/\rho)^2}ds$$
$$= \left[1-\Phi\left(-\frac{f}{\sigma}+\sigma(\lambda/\rho)\right)\right](\lambda/\rho)e^{-(\lambda/\rho)f+\frac{1}{2}\sigma^2(\lambda/\rho)^2}$$
$$= \Phi\left(\frac{f}{\sigma}-\sigma(\lambda/\rho)\right)(\lambda/\rho)e^{-(\lambda/\rho)f+\frac{1}{2}\sigma^2(\lambda/\rho)^2}. \tag{C4}$$

The third line follows from completing the square, the fourth line from the definition of the normal cdf, and the final line from the symmetry of the normal cdf.

The contribution to the log-likelihood for one observation therefore is

$$\mathcal{L}_i(\theta,\sigma,(\lambda/\rho);y_i,X_i) = \ln\left[\Phi\left(\frac{y_i-X_i\theta}{\sigma}-\sigma(\lambda/\rho)\right)\right]+\ln(\lambda/\rho)-(\lambda/\rho)(y_i-X_i\theta) + \frac{1}{2}\sigma^2(\lambda/\rho)^2, \tag{C5}$$

and the log-likelihood function for a sample of N observations is

$$\mathcal{L}(\theta,\sigma,(\lambda/\rho);\{y_i,X_i\}_{i=1}^N) = \sum_i \mathcal{L}_i(\theta,\sigma,(\lambda/\rho);y_i,X_i) \tag{C6}$$

Differentiate Equation (C6) with respect to the parameters to arrive at the moment conditions for the model:

$$\frac{\partial\mathcal{L}}{\partial\theta}: \quad \sum_i\left[\frac{\phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)\frac{-X_i'}{\hat{\sigma}}}{\Phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)}+(\widehat{\lambda/\rho})X_i'\right]=0 \tag{C7}$$

$$\frac{\partial\mathcal{L}}{\partial\sigma}: \quad \sum_i\left[\frac{\phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)\left(-\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}^2}-(\widehat{\lambda/\rho})\right)}{\Phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)}+\hat{\sigma}(\widehat{\lambda/\rho})^2\right]=0 \tag{C8}$$

$$\frac{\partial\mathcal{L}}{\partial(\lambda/\rho)}: \quad \sum_i\left[\frac{\phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)(-\hat{\sigma})}{\Phi\left(\frac{y_i-X_i\hat{\theta}}{\hat{\sigma}}-\hat{\sigma}(\widehat{\lambda/\rho})\right)}+\frac{1}{(\widehat{\lambda/\rho})}-(y_i-X_i\hat{\theta})+\hat{\sigma}^2(\widehat{\lambda/\rho})\right]=0, \tag{C9}$$

where hat denotes the maximum-likelihood estimates.

Let $\hat{f}_i$ be the empirical residuals and N the total number of observations. From the first-order condition for θ

$$\frac{1}{N}\sum_i\left[\left(-\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)\frac{1}{\hat{\sigma}}}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}+\widehat{(\lambda/\rho)}\right)X_i{}'\right]=0. \tag{C10}$$

From the first-order condition for (λ/ρ),

$$\frac{1}{N}\sum_i\left[\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}\right](-\hat{\sigma})+\frac{1}{\widehat{(\lambda/\rho)}}+\hat{\sigma}^2\widehat{(\lambda/\rho)}=\frac{1}{N}\sum_i\hat{f}_i. \tag{C11}$$

If the model includes a constant so that the first element of X_i is a constant, Equation (C10) implies

$$\frac{1}{N}\sum_i\left[\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}\right]=\hat{\sigma}\widehat{(\lambda/\rho)}, \tag{C12}$$

and plugging into Equation (C11),

$$\frac{1}{N}\sum_i\hat{f}_i=\frac{1}{\widehat{(\lambda/\rho)}}. \tag{C13}$$

Equation (C13) is the empirical analog of the moment condition

$$E[\epsilon_i+\xi_i]=\frac{1}{(\lambda/\rho)}. \tag{C14}$$

From the first-order condition for σ,

$$\frac{1}{N}\sum_i\left[\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)\left(-\frac{\hat{f}_i}{\hat{\sigma}^2}\right)}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}\right]=-\hat{\sigma}\widehat{(\lambda/\rho)}^2+\frac{1}{N}\sum_i\left[\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}\right]\widehat{(\lambda/\rho)}. \tag{C15}$$

Using the first-order condition for θ and assuming that the model includes a constant, Equation (C15) implies

$$\frac{1}{N}\sum_i\left[\frac{\phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)\hat{f}_i}{\Phi\left(\frac{\hat{f}_i}{\hat{\sigma}}-\hat{\sigma}\widehat{(\lambda/\rho)}\right)}\right]=0 \tag{C16}$$

The expected value of ϵ_i conditional on $\epsilon_i+\xi_i=f_i$ is

$$E[\epsilon_i|\epsilon_i + \xi_i = f_i] = \sigma \left[\frac{\phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right) f_i}{\Phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right)} \right]. \tag{C17}$$

The moment condition in Equation (C16) is the empirical analog to the unconditional moment condition

$$E[E[\epsilon_i|\epsilon_i + \xi_i = f_i]] = 0. \tag{C18}$$

The first-order condition for θ, Equation (C10), is the empirical analog to the moment condition

$$E\left[\left(-\frac{\phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right)\frac{1}{\sigma}}{\Phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right)} + (\lambda/\rho) \right) X_i^{'} \right] = 0 \tag{C19}$$

We can interpret $-\frac{\phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right)\frac{1}{\sigma}}{\Phi\left(\frac{f_i}{\sigma} - \sigma(\lambda/\rho)\right)} + (\lambda/\rho)$ as a generalized residual for the model, which should be orthogonal to X_i. Moment condition (C19) is the condition that f_i be distributed as the sum of a normal with standard deviation σ and an exponential with parameter (λ/ρ).

To summarize, the parameters are identified from the conditions that the unconditional mean of the residuals equals the mean of the exponential (ρ/λ), the unconditional mean of the normal part of the errors equals zero—the condition identifies σ; and finally that the density of the empirical residuals matches the density of the sum of normal and exponential random variables as closely as possible—identifying θ.

References

Aigner, D., C. A. K. Lovell, and P. Schmidt, 1977, "Formulation and Estimation of Stochastic Frontier Production Function Models," *Journal of Econometrics*, 6, 21–37.

Altunbas, Y., E. P. M. Gardener, P. Molyneux, and B. Moore, 2001, "Efficiency in European Banking," *European Economic Review*, 45, 1931–1955.

Berger, A. N., and L. J. Mester, 1997, "Inside the Black Box: What Explains Differences in the Efficiencies of Financial Institutions," *Journal of Banking and Finance*, 21, 895–947.

Chakravarty, S., and A. Sarkar, 1999, "Liquidity in U.S. Fixed Income Markets: A Comparison of the Bid-Ask Spread in Corporate, Government and Municipal Bond Markets," Working Paper, Federal Reserve Bank of New York.

Duffie, D., N. Gârleanu, and L. Pedersen, 2005, "Over-the-Counter Markets," *Econometrica,* 73, 1815–1874.

Green, R. C., 1993, "A Simple Model of the Taxable and Tax-Exempt Yield Curves," *Review of Financial Studies*, 6, 233–264.

Greene, W. H., 2002, *Econometric Analysis* (5th ed.), Prentice Hall, Upper Saddle River, New Jersey.

Härdle, W., 1990, *Applied Nonparametric Regression*, Cambridge University Press, New York.

Habib, M., and A. Ljungqvist, 2005, "Firm Value and Managerial Incentives," *Journal of Business*, 78, 2053–2094.

Harris, L., and M. Piwowar, 2006, "Municipal Bond Liquidity," *Journal of Finance*, 61, 1330–1366.

Hong, G., and A. Warga, 2000, "An Empirical Study of Bond Market Transactions," *Financial Analysts Journal*, 56, 32–46.

Hong, G., and A. Warga, 2004, "Municipal Marketability," *Journal of Fixed Income*, 14, 86–95.

Hunt-McCool, J., S. C. Koh, and B. Francis, 1996, "Testing for Deliberate Underpricing in the IPO Premarket: A Stochastic Frontier Approach," *Review of Financial Studies*, 9, 1251–1269.

Koop, G., and K. Li, 2001, "The Valuation of IPO and SEO Firms," *Journal of Empirical Finance*, 8, 375–401.

Kumbhakar, S. C., and C. A. K. Lovell, 2003, *Stochastic Frontier Analysis*, Cambridge University Press, New York.

Parkinson, M., 'The Extreme Value Method for Estimating the Variance of a Rate of Return,' Journal of Business, 53, 61–65.

Schultz, P., 2001, "Corporate Bond Trading Costs: A Peek Behind the Curtain," *Journal of Finance*, 56, 677–698.