

Financial Intermediation and the Costs of Trading in an Opaque Market*

by

Richard C. Green
Burton Hollifield
and
Norman Schürhoff

Graduate School of Industrial Administration
Carnegie Mellon University

April 20, 2004

*We would like to thank Tal Heppenstal of PriMuni LLC for providing the dataset and providing information about the structure of the muni market. Bruno Biais, David Brown, Larry Harris, and Chester Spatt provided useful comments on earlier drafts.

Abstract

Municipal bonds trade in opaque, decentralized broker-dealer markets in which price information is costly to gather. Whether dealers in such a market operate competitively is an empirical issue, but a difficult one to study because data in such markets is generally not centrally recorded. We analyze a comprehensive database of all trades between broker-dealers in municipal bonds and their customers. The data is only released to the public with a substantial lag, and thus the market was relatively opaque to the traders themselves during our sample period. We use our sample to estimate the cross-sectional determinants of the dealer markups. We find that dealers earn lower average markups on larger trades with customers, even though larger trades lead the dealers to bear more risk of losses. We formulate and estimate a simple structural bargaining model that allows us to estimate measures of dealer bargaining power and relate it to characteristics of the trades. The results suggest dealers exercise substantial market power in their trades with customers. Our measures of market power decrease in trade size and increase in variables that indicate the complexity of the trade for the dealer.

1 Introduction

The market for municipal bonds in the U.S. is a large and important financial market by any measure. Aggregate municipal bond holdings currently amount to roughly \$1.75 trillion.¹ New issues in recent years have averaged close to \$300 billion per year. Along with its size, the municipal market is also 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 ten and fifteen thousand separate municipal bond issues per year. In 2001 the average issue size was \$21 million.² Since bonds are typically issued in series, each issue in turn is likely to consist of 10-30 separate bonds with different maturities. Second, since municipals are tax-exempt, they are in large measure held by individual retail investors (35% in 2001), insurance companies (10%), or in bank and personal trusts (6.5%). Municipal bonds are therefore unlikely to trade frequently.

Third, municipals are traded in decentralized broker-dealer markets. There is no centralized exchange, and to obtain quotes a buyer or seller must call multiple dealers or solicit bids from them. The lack of a centralized exchange makes comparison shopping relatively costly. Fourth, many states only exempt their own municipalities' bonds from state income tax, further contributing to local segmentation. Finally, either for political reasons or because of local 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.

We analyze a database that allows us to observe the trading of broker-dealer intermediaries in the municipal bond market. The data has been gathered by the Municipal Securities Rulemaking Board (MSRB) as a first step in their efforts to move municipal bond trading towards a more transparent system. The database records every transaction in municipal bonds by registered broker-dealers between May 1, 2000 and July 31, 2001. A notable feature of our sample is that it contains a large number of transactions — over 8 million trades — over a longer time than in a typical microstructure transactions based sample.

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

²Data on the number of underwritings is available at <http://www.financialservicesfacts.org>.

These data are only released to the public after a time lag. During our sample period this lag was one day for bonds that traded more than four times during the day, and one month for all other trades. While the database does not identify the specific broker-dealer associated with a given trade, it 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 these spreads, including the price concessions to quantity and inventory risk faced by the intermediary.

We develop a simple model that decomposes dealer profits on a trade into a component due to the cost of facilitating the trade, a zero-mean forecast error, and a measure of the dealer's bargaining advantage or market power. We estimate the model as a Stochastic Frontier Model. The results are informative about how dealers' market power varies with the type of trade.

There is a long-standing debate in the regulatory community and among researchers concerning the costs and benefits of centralization and transparency in financial markets. 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. A significant difficulty in resolving these questions is that, almost by construction, it is difficult to obtain transactions-based data in markets that are fragmented and opaque. 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. Our analysis can contribute to a better understanding of such broad questions by measuring the risks and rewards earned by intermediaries in a fragmented market of the sort that has so far been the most difficult to study.

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 since the MSRB data has become available. Many argue that the spreads are unreasonably high. A popular web site (www.MunicipalBonds.com) lists the "Worst Ten Spreads" for trades in various categories, along with "Red Flags" and "Worst Spread Reports." MSRB rules require that trades with customers should only be executed at prices that are "fair and reasonable, taking into consideration all relevant factors." Both NASD and the SEC have examined specific cases and complaints concerning municipal bond trades. The difficulty for

researchers and policy makers lies in identifying the relevant factors and quantifying when variation attributable to those factors is fair and reasonable. Absent a natural experiment, this requires a structural model to disentangle competitive compensation for the costs of intermediating trades from the exercise of market power. The SEC ruled that markups on trades ranging from 1.42% to 5% “substantially exceeded accepted industry practice,” while 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 costs of financial intermediation, will necessarily be conservative and identify only extreme examples of “unfair and unreasonable” markups.

We use our sample to answer the following questions. How much does it cost to trade in the municipal bond market? How variable are these costs, and what are the determinants of this variability? How are the costs affected by liquidity and seasoning of the issue? How much do dealers earn for matching buyers and sellers, and how does the revenue depend on the risk that the dealers appear to bear or effort they expend? To what extent is variation in the costs to investors evidence of market power for dealers?

Since we cannot directly link individual trades to specific dealers, we must infer their trades and profits indirectly. For example, we study trades that appear to be direct exchanges of bonds through a dealer. The same par value of the same bond is purchased and then sold in a very short time period. The costs of such trades for customers represent a measure of the rewards to pure intermediation—compensation for the search costs of identifying counter-parties—along with any rents that accrue to the intermediary through market power. Alternatively, we consider more generous filters, such as purchases by a dealer followed by a sequence of sales to customers or interdealer trades that return dealer inventories to their original level. Because the individual bonds trade infrequently, we can be reasonably confident that these trades are associated with the same original block of bonds. Our empirical findings are very robust to the method for isolating specific round-trip transactions.

Empirical methods alone are of limited use in identifying that portion of dealer profits due to market power. We develop a very simple theoretical and empirical model of the interaction

³See “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.

between dealers and their customers in which the expected profits to the dealer reflect the dealer's costs and 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 stochastic frontier. The stochastic frontier represents the expected efficient markup the customer would obtain if dealers were always driven to their reservation value, as they would be if the provision of dealer services were perfectly competitive. The observed markup differs from dealer 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 bargaining power.

Our empirical results suggest dealer costs depend on factors one would expect influence the costs and risk of intermediating trades, such as liquidity measures for the bond and market, the size of the trade, and interest rate conditions. 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. It is also higher, for the dealers as a group, in situations where more extensive intermediation is required, such as when the trade moves through other dealers, spends more time in inventory, or is broken up into smaller blocks.

The lack of centralized recording and preservation of price information has 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 study a sample of trades executed on behalf of large insurance companies. Their data identify the institutions and dealers participating in the trade, and whether it was a buy or a sell. The authors use various techniques to estimate the bid-ask spread implicit in the observed trades, and then relate the measured trading costs to characteristics of the bonds and the current market conditions.

Schultz (2001), who studies corporate bonds, shows that trading costs decrease with the size of the trade, and are lower for institutions and dealers that are more active in the market. Hong and Warga (2000) also study corporate bonds, and compare estimates of realized spreads in the dealership market to those on the NYSE's Automated Bond System, and to monthly quotations provided by Lehman Brothers. They show spreads depend on bond age, issue size, volatility, maturity, and rating. Chakravarty and Sarkar (1999), who study corporate, government, and

municipal bonds, find that trading costs decrease as volume rises and increase in the age of the bond.

Our dataset differs from that used in the previous studies in several respects. First, it is comprehensive. It includes all trades by registered broker-dealers in municipal securities, rather than those trades initiated by a subset of institutional investors. We have data on over 8 million transactions in over 500 thousand bonds issued by over 37 thousand entities. Our data do not identify dealers or buyers specifically. They only tell us whether a trade was a buy from a customer, a sell to a customer, or a trade between dealers. Thus, we can measure the aggregate flow of bonds to and from the dealers as a group, and the turnover of bonds within that group, but we cannot relate trading costs to the individual characteristics of the buyer or seller. The earlier studies mentioned above use other data sources to identify the characteristics of the particular bonds. The information in our data about the characteristics of the bonds is limited, and given the huge number of different bonds we have been unable to gather additional data on specific attributes of the bonds for the whole sample.

The paper closest to our work is Harris and Piwowar (2004), who also use the MSRB data over a similar sample period to estimate trading costs for municipal bonds. Their approach is very different from ours, reflecting differences in goals. The two papers share an interest in measuring how large the costs of trading in this market are, and understanding how these costs vary with the size and other characteristics of the trade. They also share a central difficulty: while there are many trades in many bonds, trading in any one bond is typically very infrequent. Harris and Piwowar (2004) estimate 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 analyze how the estimated costs vary cross-sectionally for different types of bonds. We focus on pairs of trades that can reasonably be assumed to represent two sides of a single intermediated transaction, and then 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 radically different approaches produce roughly similar estimates of average costs of trad-

ing, and both make clear that smaller trades are much more costly than large ones. The Harris and Piwowar (2004) results provide a richer description of how the costs vary with bond characteristics such as credit quality. That paper is purely empirical in its goals, however. Our paper uses a theoretical model to seek evidence that the high costs of trading are due to dealer 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 the paper is organized as follows. The next section describes the institutional setting and our sample. In Section 3 we study dealer profits on round-trip transactions, and illustrate their dependence on size and other characteristics of the trade. Section 4 develops a structural model that allows us to decomposed the profits into portions attributable to dealer costs and market power. Section 5 concludes. The Appendix contains detailed descriptions of procedures we used to filter the data and variable definitions.

2 The Municipal Bond Market and the Sample

From the Municipal Securities Rulemaking Board (MSRB) we have transactions data reporting all trades carried out by broker-dealers from May 1, 2000 through July 31, 2001. 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 transactions prices publicly available to market participants. The sample is self-reported, and there are many systematic errors made in entering the data, especially early in the sample period when member firms were still becoming accustomed to the system. In the appendix we describe in detail the procedures we use to identify and filter out data errors.

The transaction specific information provided about each trade include the price, the date and time of the trade, the par value traded, whether the dealer bought or sold, and whether the trade was with a customer or another dealer. There are fields for the yield at which the trade took place, which is not required but is sometimes reported. Information is also provided about the bond being traded including: the CUSIP number, a text description of the bond, the date interest begins accruing after issuance, 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 above characteristics of the bond. For example, we have been able to calculate a duration measure for each of over 8 million trades in the sample.

Table 1 provides descriptive information on the MSRB data. The first panel includes all transactions. The second panel describes trades occurring in the first ninety days after issuance for bonds that are issued during our sample period.⁴ 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 then distributed to customers primarily during the first ninety days. The dollar value of trades in newly issued bonds is \$772,401 million, only 23% of the total value of all trades. The difference between sales to customers and buys from customers for newly issued bonds, \$291,281 million, is 60% of the corresponding difference for all trades, even though transactions in new issues are less than a quarter of all transactions.

Sales to customers are approximately equal to buys from customers in seasoned issues, providing evidence of the intermediating services provided by the broker-dealers. The broker-dealers tend to buy bonds from customers in larger par amounts, and then sell the bonds in smaller blocks. The average purchase from a customer is close to five times larger than the average sale to a customer for new issues, and a bit less than three times the corresponding transaction size for seasoned issues. Inter-dealer trades are of intermediate size.

Figure 1 illustrates how volume of trade evolves over our sample period. The top line is sales to customers of seasoned issues and the bottom line is the aggregate order imbalance—the difference between sales to and purchases from customers. The series show peaks and valleys, but no broad trends over the period.

Municipal bonds are actively traded in a “when issued” market, and during the period 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 percent of the outstanding issues did not trade at all. Such patterns are very apparent in Table 2

⁴For 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, since our data suggests that transaction volume has not yet flattened out one month after issuance. Substantial transaction volume in new issues starts as early as 20 days before the dated date.

and Figure 2.

The first column in the body of Table 2 reports the unconditional average number of trades per month for the bonds that appear in our sample, and thus trade at some point during our sample period. The table reports the average number of trades as a function of the bonds' seasoning. Here, seasoning is defined as the time since issuance as a percentage of the original maturity of the bond. For example, if the bond has a one-year maturity, a trade six months after the bond is issued would have 50% seasoning. The rows of the table therefore represent deciles of a bond's life.

On average, a bond has approximately 6 trades in the first 10% of its life. But our sample only includes municipal bonds that trade—there are many municipal bonds outstanding that do not appear in our transactions-based sample because they do not trade during the sample period. The average municipal bond therefore has less than 6 trades during the first 10% of its life. By the time a bond is 10% of the way to maturity, average volume has dropped to 0.16 trades per month, and by the time it is half way to maturity there are fewer than 0.1 trades in a month.

The second column Table 2 computes the average number of trades, conditional on the bond trading at some point in a particular decile of its seasoning. Conditioning in this way gives an accurate picture of the volume in a bond, given that a transaction has taken place. Obviously, the more actively traded bonds are more likely to appear in a transaction-based sample. Still, on average, there are fewer than two trades a month once the bond is 10% of the way to maturity, and less than one when the bond is reasonably seasoned.

In Figure 2 we plot the empirical probability of trade as a function of the age of the bond. Four to six months from issuance, less than ten percent of the bonds in our sample trade at all. The probability then rises somewhat so that by two years from issuance, roughly twenty percent of the bonds in the sample are trading at least once during a month.

Figure 3 shows 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 per day stays close to two, suggesting that trades through a broker dealer in a given bond are typically limited to the intermediation of a specific transaction. Keeping in mind that customers are unlikely to even attempt transactions in more illiquid bonds, the finding suggests that dealers taking one side of a trade to provide liquidity will face substantial uncertainty and costs of search to identify a counter

party.

3 The Magnitude and Characteristics of Dealer Markups

Our data does not identify matched buy-sell transactions for specific bonds, or the specific dealers carrying out those transactions. We simply 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 leave little doubt that they are two ends of a transaction intermediated by a specific dealer. We focus on two types of these transactions. 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, and 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. Such transactions are therefore likely to involve counter-parties that have been identified beforehand.

Of course, some of the immediate matches may involve counter-parties that were not identified by the intermediary prior to initiating the first transaction. Indeed, there may be cases where the initial purchase is by one intermediary, the subsequent sale is by another, and it is purely coincidence that the par value of the trades is identical. The information from the previous section about the frequency of trade, however, makes such events seem quite unlikely, particularly as the time between the two trades becomes small. As is apparent from Figure 3, 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.

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. 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 these transaction represent rewards to the dealers as a group for the intermediation they provide. We refer to such transactions as “round-trip” transactions. The appendix provides the details about how we match these round-trip transactions.

A total of 908,638 purchase transactions in seasoned bonds, 51.5% of all purchases from a customer, can be separated out as round-trip transactions. In 79% of all cases, the bonds purchased from the customer are resold in one piece. In the remaining cases, the bonds are broken up and resold in smaller quantities. The round-trip transactions exclude all transactions where dealers do the reverse of a purchase from a customer followed by sales to customer: i.e., bundling and reselling the bonds as a larger package. When the bond is split up, the average number of customer sales associated with a single customer purchase is 3.4. In 21% of the transactions, the initiating dealer sells at least a portion of the initial bonds to another dealer. In exceptional cases, there can be up to 65 intermediate dealer transactions until the whole bond lot reaches retail customers. Conditional on the initial dealer passing the bonds to another dealer, there are on average 1.95 inter-dealer trades before the aggregate dealer inventory of that bond returns to its initial level.

Table 3 summarizes the markups the dealers earn on the round-trip transactions, measured in basis points, and sorted by trade size. The average markup on a round-trip transaction is 2.1%. The most striking and robust finding is that the percentage markups decrease with trade size, independently of how the trade is executed. The relationship between size and the risk or difficulty of intermediating the transactions, however, is not a simple one.

First, it is evident that dealers do bear some risk in their role as intermediaries. The third column in Table 4 shows the overall frequency of a loss is 1.4%. A total of 6.1% of the transactions are transacted with no markup. Dealer profits are lower on transactions that are immediate matches than for transactions that take longer to match. Markups on immediate matches are reported in the second column in Table 3. The average markup for immediate matches is 1.6%, with a loss-probability of approximately 1.5%. The fourth and fifth columns in Table 3 show that when the bonds remain in dealer inventories for more than a day, dealers as a group bear more risk and earn higher returns.

Dealers also appear to earn higher compensation for transactions that involve larger numbers of trades or participants. Profits to dealers as a group are higher for trades that require more extensive or elaborate intermediation. The sixth and seventh columns of the table show that dealers earn more on trades that must be broken up to be sold in pieces rather than a single block. The general patterns appear reasonably robust across trade sizes in Table 3. These trades also appear riskier

in that for every trade size except in the \$10,000 to \$25,000 range, the standard deviation of the realized markup is higher when there are larger numbers of trades.

The pattern of dealer markups across trades size and type suggests search frictions force dealers to trade off risk and return. Selling to another dealer may be a costly option. Whenever no direct retail counter party is available, dealers may unwind by selling to another dealer at lower markups or losses. Twenty-one percent of the time the initial dealer sells some or all of the initial purchase from a customer to another dealer. The markups earned by the dealers as a group from such transactions are 2.2% on average, and the standard deviation is 1.8%. It is unclear whether this is because other dealers simply drive a harder bargain, or because they are willing to provide liquidity in falling markets when customers are not buying.

Nevertheless, there is certainly some evidence bargaining power is tilted towards larger customers. While markups decline with trade size, risk does not. Table 4 contains several statistics describing the different size categories. The second column reports the number of trades. The largest numbers of trades are in the small to intermediate size category. The most profitable trades for dealers, however, are larger, in the range from \$500 thousand to \$750 thousand.

The largest trades, however, are not particularly profitable, consistent with the extremely low margins evident for large transactions in Table 3. Profits on the largest trades are low on average, in both dollar and percentage terms, despite evidence that the largest trades are the most risky. The third column in the Table 4 reports the frequency with which dealer profits are negative. The loss frequencies are low and relatively stable across all trade sizes, except for the largest trades. The loss frequencies for largest trade sizes increase dramatically.

The loss frequencies for small and intermediate sized trades are relatively low, even though the standard deviation of the markups in Table 3 are relatively large. The loss frequency is not monotonically related to the standard deviation of the markups because the distribution of markups is highly positively skewed.

Figure 4 shows the frequency distributions for four different trade-size categories. 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.⁵ It is obvious from

⁵Härdle (1990) provides a textbook introduction to kernel smoothing methods.

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

Figure 4 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), while very few trades occur at spreads in excess of 8%. Except for the very largest trades, there is considerable probability mass above 2%. This suggests that, 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.

4 Dealer Market Power and Transaction Size

In the previous section we established that the markups dealers earn on round-trip transactions 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?

In this section we speak to this question by developing a simple model that decomposes the observed markup into components associated with the cost to the dealer and the dealer’s relative bargaining power in dealing with the customer. We then estimate the model as a stochastic frontier model.

Consider transactions between a dealer, indexed by i , and a seller, indexed by j . We assume that on coming to a dealer the seller arrives with a reservation value v_j , which represents his outside opportunities. The reservation value can be viewed, for example, as the continuation value

associated with searching for a different dealer. 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_i , and its expectation as $E(p_i|X)$, where X is a set of conditioning variables, observable to the dealer and the seller. The anticipated cost of intermediating the trade is $c_i(X, \theta)$, where θ is a set of parameters to be estimated. Finally, let p^* be the price the dealer offers to the seller.

We assume that the dealer is risk-neutral with indirect utility function equal to the expected profit from purchasing the bond today for p^* , reselling it in the future for an expected price of $E(p_i|X)$, and paying the expected intermediation cost, $c_i(X, \theta)$:

$$U_i(p^*; X) = E(p_i|X) - c_i(X, \theta) - p^*. \quad (1)$$

The seller is also assumed to be risk-neutral with indirect utility function equal to the price she receives for selling the bond today, p^* , less her reservation value for the bond, v_j :

$$U_j(p^*; X) = p^* - v_j. \quad (2)$$

A price p^* such that both indirect utilities are positive only exists when there are gains to trade. If there are no gains to trade the game ends, and no transaction takes place. If there are gains to trade, seller and dealer engage in an alternating offer game with possibility of breakdown, the solution of which can be described by the generalized Nash solution.

Let ρ_{ij} be the bargaining power of dealer i relative to that of seller j , satisfying $\rho_{ij} \in [0, 1]$. If $\rho_{ij} = 0$, the seller has all the bargaining power and if $\rho_{ij} = 1$, the dealer has all the bargaining power.

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

$$\max_{p^*} (E(p_i|X) - c_i(X, \theta) - p^*)^{\rho_{ij}} (p^* - v_j)^{1-\rho_{ij}} \quad (3)$$

subject to the participation constraints

$$E(p_i|X) - c_i(X, \theta) - p^* \geq 0, \quad (4)$$

$$p^* - v_j \geq 0. \quad (5)$$

Condition (4) requires the price the dealer pays must be less than the expected net revenues from reselling the bond, and condition (5) requires the price the seller gets for the bond exceeds her reservation value.

The first-order condition in the interior is:

$$-\rho_{ij}(E(p_i|X) - c_i(X, \theta) - p^*) + (1 - \rho_{ij})(p^* - v_j) = 0. \quad (6)$$

Solving (6) for p^* , the equilibrium offer price is

$$p^* = \rho_{ij}v_j + (1 - \rho_{ij})(E(p_i|X) - c_i(X, \theta)). \quad (7)$$

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 ($\rho_{ij} = 0$), the offer price is equal to the expected cash flows that the dealer receives from reselling the bond, $E(p_i|X) - c_i(X, \theta)$, and the dealer makes zero expected profits. If the dealer has all the bargaining power ($\rho_{ij} = 1$), 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 he will in turn sell them for p_i , which will differ from the expected price by an expectational error with mean zero:

$$\epsilon_i \equiv p_i - E(p_i|X). \quad (8)$$

The realized markup will consist of the expectational error, along with that portion of the surplus

the dealer is able to extract from the seller, which we will denote u_{ij} . To see this, note that

$$\begin{aligned} p_i - p^* &= E(p_i|X) - p^* + \epsilon_i \\ &= c_i(X, \theta) + \epsilon_i + u_{ij} \end{aligned} \tag{9}$$

where

$$\begin{aligned} u_{ij} &\equiv E(p_i|X) - c_i(X, \theta) - p^* \\ &= \rho_{ij}(E(p_i|X) - c_i(X, \theta) - v_j). \end{aligned} \tag{10}$$

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

$$u_{ij} \geq 0. \tag{11}$$

Condition (11) follows because $\rho_{ij} \in [0, 1]$ and because the participation constraints, (4) and (5) can both be satisfied if and only if the gains to trade are positive: $E(p_i|X) - c_i(X, \theta) - v_j \geq 0$.

The model in equation (9) fits naturally into the econometric specifications that can be estimated using Stochastic Frontier Analysis. Greene (2002) and Kumbhakar and Lovell (2003) provide 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 to study IPO underpricing by Hunt-McCool, Koh and Francis (1996) and Koop and Li (2001), efficiency in the banking industry by Berger and Mester (1997) and Altunbas, Gardener, Molyneux and Moore (2001), and the role of incentives in mitigating agency costs by Habib and Ljungqvist (2003).

In our application, the cost of intermediating the trade to the dealer, $c_i(X, \theta)$, can be viewed as the “efficient” markup from the standpoint of the seller. 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, observed markups would vary around this cost by the expectational error, ϵ_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 u_{ij} that reflects the dealer’s market or bargaining power and the seller’s reservation price. Measures of the relative “size” of the errors, ϵ_i and u_{ij} provide information about the relative importance of the dealers’ market power.

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_j . We assume $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$. Reservation values are drawn from a distribution that is centered on the dealer’s expected net proceeds from selling the bond. That is,

$$E(v_j|X) = E(p_i|X) - c_i(X, \theta). \quad (12)$$

It is natural to assume potential sellers’ reservation values are “rational” in this sense. They reflect correct expectations on average. Define ν_j as the deviation of v_j around its mean

$$\nu_j \equiv E(v_j|X) - v_j. \quad (13)$$

Substituting equation (13) into equation (10),

$$u_{ij} = \rho_{ij}\nu_j. \quad (14)$$

We parameterize ν_j as a double exponential, where the density on each side of the mean is 0.5 times the exponential density with parameter λ_j . Accordingly, the distribution conditional on $\nu_j \geq 0$ is exponential with parameter λ_j . The one-sided error u_{ij} , conditional on $u_{ij} \geq 0$, therefore is exponentially distributed with parameter λ_j/ρ_{ij} .⁶ The first two moments of u_{ij} are

$$E(u_{ij}|u_{ij} \geq 0) = \frac{\rho_{ij}}{\lambda_j}, \quad (15)$$

$$\text{Var}(u_{ij}|u_{ij} \geq 0) = \left(\frac{\rho_{ij}}{\lambda_j}\right)^2. \quad (16)$$

⁶We have also experimented with using a truncated normal distribution for the one-sided error. Our empirical findings are robust to this change.

The distribution of the asymmetric error u_{ij} can depend on observable variables. 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 $\frac{\rho_{ij}}{\lambda_j}$ that is a log-linear function of the explanatory variables $Z = (Z_k)_{k=1, \dots, K}$:

$$\frac{\rho_{ij}}{\lambda_j} = a_0 \prod_{k=1}^K e^{a_k Z_k}. \quad (17)$$

The variables Z_1, \dots, Z_K 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:

$$\sigma_i = b_0 \prod_{k=1}^K e^{b_k Z_k}. \quad (18)$$

If $a_0 = 0$, then $\frac{\rho_{ij}}{\lambda} = 0$. In this case, the sellers have all the bargaining power and there is no asymmetric error in the realized markup. If $a_0 \neq 0$, then $\frac{\rho_{ij}}{\lambda} \neq 0$. In this case, the dealers have bargaining power and there is an asymmetric error in the realized markup. Variables that are positively related to $\frac{\rho_{ij}}{\lambda_j}$ in this specification suggest higher bargaining power for the dealers.

We model the heterogeneity across customers through variation in ρ rather than in the distribution of the customer reservation prices. There is no natural econometric or economic way to distinguish between variation in ρ or variation in λ across customers, however. One would expect more sophisticated customers to come to the market with a better understanding of the dealers' costs and resale opportunities. More sophisticated customers would therefore demand prices leading to lower dealer profits. This could be interpreted either as a more concentrated distribution for their reservation prices, in which case we could allow λ to depend on characteristics of the trade,

or as greater bargaining power for the customer and thus heterogeneity in ρ .

Finally, we parameterize the cost function $c_i(X, \theta)$ as a linear function of the explanatory variables:

$$c_i(X, \theta) = \theta_0 + \sum_{l=1}^L \theta_l X_l \quad (19)$$

where L is the number of regressors.

We report estimates of the stochastic frontier models for both narrowly and more broadly defined samples. Because the data does not directly identify the dealers involved in trades, we must focus on trades where purchases are followed by sales of the same par value in the same bond. A narrowly defined sample of these trades, where the two sides of the trade are close in time—“immediate matches,” may be less representative of the dealers’ routine activities. On the other hand, a broader sample is more likely to exhibit selection bias of a different sort. If the trades are separated in time, and our selection criteria requires that there are no additional purchases (potentially by other dealers) between the buys and sells, we are biasing our sample away from more actively traded bonds. The general conclusions of our analysis are very robust to the selection criteria we employ, however.

Table 6 reports the results of fitting the empirical model in equations (17)–(19) to a narrowly defined sample conditioning on the variables reported in Table 5. The sample consists of immediate matches only—seasoned bonds (more than 90 days old) and transactions where the purchase from a customer was sold off to customers on the same day with no intermediate purchases. There are 221,032 observations used in the estimation.

In order to aid the interpretation of the coefficients, we transformed the independent variables by subtracting their broad-sample means. The appendix provides more detailed definitions of the variables and Table 5 reports the means and standard deviations of the variables. The descriptive statistics are similar in all samples, which is reassuring there is no mechanical sample selection bias. The typical transaction size in the broad sample is \$35,000 par and the typical bond is six and one half years old, has eleven and one half years to maturity and a 4.6% coupon. The bond also trades on average 25 times during the sample period—the number of trading days is 316.

The first two columns of Table 6 report the results of fitting the model with all the parameters in the one-sided error constrained to zero. In this parameterization, the model is constrained so

that the sellers have all the bargaining power. The model is fit using maximum likelihood. The first panel in Table 6 reports the coefficients on the “efficient markup,” $c_i(X, \theta)$, the second panel of the table reports the coefficients in the residual standard deviation, and the likelihood function is reported in the final panel of the table.

Almost all the explanatory variables have significant t-statistics. 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. In environments with high interest rates, and more volatile rates, markups are larger. Types of bonds that are notoriously illiquid, such as pre-refunded bonds, earn higher markups. Bonds with higher liquidity, measured by purchase frequency or issue size (proxied by the maximum par size traded) earn dealers lower spreads.

The second two columns of Table 6 report maximum-likelihood estimates of a stochastic frontier model with specification (17)-(19), with the parameters unrestricted. The first panel of Table 6 reports the coefficients in the linear model for the “efficient” markup, or $c_i(X, \theta)$. The second and third panels report the estimates for the two error distributions. The final panel reports the value of the likelihood function.

Since the explanatory variables are demeaned in the regression, the intercept terms are estimates of the expected efficient cost and the expected loss to dealer market power on an average transaction. The constant in the estimates for the asymmetric error (123.35) swamps the constant in the cost function (2.08). The constant in the cost function has a natural interpretation of the average efficient cost of search or intermediation. Thus, much of the spread on the average transaction can be attributed to the dealer’s advantage in bargaining power.

Many of the coefficients in the specification with a one-sided error decrease relative to the model with only a symmetric error distribution, both in economic and statistical significance. Notably, the coefficient on the size of the transaction reverses sign and is of marginal significance. The portion of the markup due to bargaining power, however, is significantly negatively related to size. The reduction in the significance of the coefficients is particularly evident for liquidity related variables. For example, the dummy for pre-refunded bonds is no longer significant, and the sign on the purchase frequency variable changes. That the effects of these variables are 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 illiquid bonds rather than their effect on the dealer's costs.

Table 7 confirms that the qualitative results are robust to the sample selection. The table reports estimates using the broadly defined sample, which includes the markups for all seasoned bonds in our sample. There are 908,638 observations used in the estimation. The first two columns report estimates and t-statistics from the restricted model and the last two columns results for the unrestricted model. The estimates from the restricted model reported in the first two columns are similar to those reported in the first two columns of Table 6. The estimates reported in the third column confirm that actual costs are economically insignificant for the average transaction, but that discriminatory markups are large.

The models reported in Table 7 include dummy variables for the intermediation services interacted with the transaction size. The estimated coefficients in the cost function suggest transactions that involve other intermediaries, or involve selling the bonds to multiple buyers, earn higher markups. This is what one would expect, since the dealers are clearly providing more service in these situations. The dummy variables are interacted with many of the other variables. Allowing for such interaction terms is both empirically and economically important, since trade size proxies for heterogeneity in customer type. For instance, one would expect that pre-negotiation is more pervasive for larger trades. Cases of large purchases remaining in dealer inventory over night suggest that the seller demanded liquidity. Such an interpretation is confirmed in Table 7, by the positive coefficients on *Overnight* and *Days In Inventory* \times *Overnight* for par greater than \$100,000 in the specification of the asymmetric error.

The estimates of the model for the asymmetric error term also show that interdealer involvement in large transactions triples the asymmetric component of the markup, while interdealer involvement in small transactions reduces the asymmetric component by a third. The coefficient of 1.11 on *No Split, Interdealer* means the expected one-sided error is multiplied by $\exp(1.11) = 3.03$ if the dummy equals one for par value in the range $[0.5, \infty)$. Similarly, the coefficient 0.29 means the expected markup is multiplied by $\exp(0.29) = 1.336$ if the dummy is one for par values in $(0.0, 0.1)$. The large coefficients on *Overnight* in the cost function, and the fact that they are decreasing with par, indicate there are substantial fixed costs to holding inventory, and that search costs, as reflected in

the intercept for the cost function, are in comparison small.

Variables for how the trade is handled, such as whether it is split and whether other dealers are involved, appear significant in both the efficient cost function and in the specification of the asymmetric error. The positive coefficients for smaller trades suggest they increase the cost of processing these trades to the dealers. One would expect this as more intermediation is clearly involved in such trades. These variables also appear as significant in conditioning the distribution of the one-sided error term. Trades that involve more extensive intermediation earn higher markups than the apparent increase in the dealer's cost. Dealers appear to undertake these transactions only in situations where they expect to earn high returns.

A partial explanation for the result that splitting tends to increase 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 in order 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 they will probably need to sell the bonds to another dealer, and appear to do so only when the terms are particularly advantageous.

Finally, the role of the time the bonds spend in inventory, as measured by *Days In Inventory* is more complex, as one might expect given that bonds remain in inventory only when dealers fail to find an attractive price quickly. Longer time in inventory reduces the efficient cost, but slightly increases the asymmetric component of the spread. Dealers appear to accept trades that they know will require time in inventory only when the cost to them is low, and their bargaining power with the customer is high.

5 Conclusion

Corporate and municipal bonds have traditionally been traded primarily in opaque, decentralized markets in which price information is costly to gather. This 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. The extent to which they operate competitively is an empirical issue. Data on trade in these markets, however, has been difficult to obtain in the past, and even when such data is 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 is only released to the public with a substantial lag, and thus the market remained relatively opaque to the traders themselves even though the data is available to researchers. Because the bonds trade infrequently, this transactions data is low frequency for each bond, but because there are so many bonds we have hundreds of thousands of observations.

We show the pattern of dealer profits in round-trip transactions implies they earn lower average profits on larger trades, even while they bear more risk. We formulate a simple structural model that allows us to estimate measures of dealer bargaining power and relate it to characteristics of the trades. The model can be estimated as a Stochastic Frontier Analysis model. The results suggest 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.

References

- Aigner, D., Lovell, C.A.K., and Schmidt, P., 1977, "Formulation and Estimation of Stochastic Frontier Production Function Models," *Journal of Econometrics* 6, 21-37.
- Altunbas, Y., Gardener, E.P.M., Molyneux, P., Moore, B., 2001, "Efficiency in European Banking," *European Economic Review* 45, 1931-1955.
- Berger, A.N., Mester, L.J., 1997, "Inside the Black Box: What Explains Differences in the Efficiencies of Financial Institutions," *Journal of Banking and Finance* 21, 895-947.
- Chakravarty, Sugato and Asani 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.
- Greene, William H., 2002, *Econometric Analysis*, Prentice Hall, 5th Edition.
- Härdle, W., 1990, *Applied Nonparametric Regression*, Cambridge University Press.
- Habib, M., Ljungqvist, A., 2003, "Firm Value and Managerial Incentives," working paper, NYU.
- Harris, L., Piwowar, M., 2004, "Municipal Bond Liquidity," working paper, Securities and Exchange Commission.
- Hong, Gwangheon and Arthur Warga, 2000, "An Empirical Study of Bond Market Transactions," *Financial Analysts Journal*, 56,
- Hunt-McCool, J., Koh, S.C., Francis, B., 1996, "Testing for Deliberate Underpricing in the IPO Premarket: A Stochastic Frontier Approach," *Review of Financial Studies* 9, 1251-1269.
- Kumbhakar, Subal C. and C. A. Knox Lovell, 2003, *Stochastic Frontier Analysis*, Cambridge University Press.
- Koop, Gary and Kai Li, 2001, "The Valuation of IPO and SEO Firms," *Journal of Empirical Finance* 8, 375-401.
- Schultz, Paul, 2001, "Corporate Bond Trading Costs: A Peek Behind the Curtain," *Journal of Finance*, 56, 677-698.

A Appendix

A.1 Data Filter

A number of transactions in the MSRB database are missing specific descriptions of the bond except its CUSIP number, and either the transaction price or the yield are not reported. MSRB reporting rules do not explicitly require broker/dealers to report both transaction price and yield.

In the case of missing bond description items such as coupon rate, issue date or maturity date, we perform searches over the entire database to recover the missing data. In many cases this is successful.

In case transaction price or yield information is missing we use a log-linear regression approach to recover the price-yield relationship using all observations in the database on the given issue. We then use reported yields to predict transaction prices or the other way around, depending on which information is missing.

More specifically, we run three sets of ordinary-least-squares regressions of log-price on yield to obtain accurate coefficient estimates. For robustness we first mechanically eliminate outliers. We restrict attention to price and yield observations within two standard deviations from their corresponding sample median. MSRB reporting rules for yields differ between transactions effected above or below par (see MSRB Rule G-12 (c) (v) (I) for details). We therefore perform two regressions in which we only include transaction prices above or, respectively, below par, and a third using all observations. Depending on whether the reported yield exceeds the coupon rate, coefficient estimates from the first or the second regression are used. Coefficient estimates from the third regression are used whenever in either of the former regressions the R-squared is below 33%, or the number of observations in the sub-samples are less than six and less than one third of all observations. Given the small number of observations for many bond issues, we found these criteria effective at producing sensible values. Finally, if the R-squared in the overall regression is less than 33%, or the total number of observations is insufficient to apply OLS in the first place, we recover the transaction price from the reported yield assuming that the yield is reported as yield-to-maturity.

A.2 Sample Selection

The MSRB database identifies neither the broker/dealer intermediating the transaction nor the retail or institutional buyer/seller. The MSRB database, however, allows to identify the type of transaction using an identifier (purchase from customer, sale to customer, inter-dealer transaction) and the time of the transaction up to a minute. We therefore use a matching algorithm to identify purchase and sales tickets that are likely to be associated with the same dealer.

We match up dealer purchases with consecutive sales to customers up until the next purchase from a customer in the database. If the par amount of a sales ticket is higher than the corresponding purchase amount, the sales transaction is ignored. All other sales transactions are added up. If the cumulative amount of these sales before the next purchase equals the par of the preceding purchase, we have a match. If there is a one-to-one match between a purchase and a sale, we say that these are no-splits. Otherwise the purchase is split into several smaller units. We record the time of the purchase and the (par weighted-average) time of the sale(s). In addition we determine whether there was a transaction between dealers in that issue at any time between the corresponding purchase and sale dates.

Broad Sample All matched transaction pairs in seasoned bonds (> 90 days since issuance).

Narrow Sample In addition to the requirements of the broad sample, we eliminate matched pairs where the purchased bond is split into smaller lots, pairs with intervening inter-dealer transactions, and pairs where the day of the sale is different from the day of the purchase.

A.3 Definitions of the Explanatory Variables

Markup Difference between the natural logarithm of the purchase and the resale price; winsorized at 0.5% and 99.5% and measured in basis points.

ln(Par) Natural log of the par value of the transaction. Measured in millions of dollars.

Maturity Years between date of transaction and maturity of bond.

Coupon Coupon rate of the bond. Measured in per cent.

Zero Coupon Bond Indicator variable for an original issue pure discount bond.

Age Age of the bond on the day of the purchase.

Transaction Frequency The total number of purchase transactions from customers per business day in the bond over our sample period. The number of business days during the sample period is 316.

Issue Size Natural log of the par value of the largest trade in the bond over our sample period. This provides information about the size of the issue.

Volume The total value of trades in bonds from that state on the day of the transaction. Municipal markets are segmented by state due to differences in state tax treatment and state income tax rates. This is a measure of the liquidity of bonds from a particular state on a given day, and these bonds are close substitutes for the bond being traded. Measured in billions of dollars.

Order Imbalance Purchases from customers less sales to customers on that day in seasoned bonds issued by municipalities in that state. Measured in billions of dollars.

New Issues Sales of new bonds in the same state on the same day. Measured in millions. New issues are more actively traded, so the frequency of trades in issues that are close substitutes for the bond in question may affect customers' reservation prices. Measured in billions of dollars.

Treasury Yield Closing yield of the maturity matched treasury bond on the same day as the purchase transaction. Measured in per cent.

Treasury Range Intraday range of the maturity matched treasury bond on the same day as the purchase transaction. Measured in basis points. This measures the volatility of interest rates on the day in question.

Overnight Bond stays in the dealer's inventory for at least one business day.

Days In Inventory Days the bond stays in dealer inventory.

No Splitting, Interdealer Dummy variable for round-trips where the bonds are sold in a block, but 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.

Revenue Bond Revenue bond.

Tax Revenue Bond Bond backed directly by tax revenues from a specific source.

General Obligation Bond General obligation bond.

Guaranteed Bond is guaranteed by a municipal bond insurance agency.

Prerefunded Bond has been prerefunded. These bonds are backed by trusts invested in treasuries, and are therefore riskless.

Certificates of Partnership Certificate of Partnership.

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.

Utilities Construction Bond issued to finance utilities construction.

Redevelopment Issuer is a redevelopment authority.

Improvement Bond issued for facilities improvements.

Facilities Construction Bond issued to finance facilities construction.

Refinancing Bond used to refinance another issue.

Table 1: Descriptive statistics for MSRB transactions data, 5/1/2000 to 7/31/2001

	All Trades	Sales to Customers	Purchases from Customers	Interdealer Trades
All Issues				
Transactions	8,330,570	4,804,997	1,864,374	1,661,199
Issues	527,440	514,942	391,268	283,222
Issuers	37,403	36,979	34,899	28,254
Volume (Million \$)	3,228,549	1,639,896	1,154,006	434,646
Transaction Size (\$)	388,199	342,242	619,133	261,646
Coupon (%)	5.40	5.33	5.55	5.42
Age	4.26	3.58	5.92	4.36
Maturity	14.43	15.3	13.26	13.22
Modified Duration	9.57	9.8	8.51	10.06
New Issues				
Transactions	1,860,204	1,343,045	93,362	423,797
Issues	164,162	161,354	42,926	93,834
Issuers	11,779	11,711	6,967	7,696
Volume (Million \$)	772,401	444,571	153,290	174,540
Transaction Size (\$)	418,316	334,326	1,650,016	411,848
Coupon (%)	5.08	5.11	5.15	4.96
Age	0.04	0.03	0.05	0.03
Maturity	14.46	14.98	14.33	12.84
Modified Duration	9.71	9.73	9.18	9.76
Seasoned Issues				
Transactions	6,441,568	3,438,683	1,766,079	1,236,806
Issues	367,840	357,684	352,841	197,188
Issuers	34,035	33,574	33,457	26,221
Volume (Million \$)	2,280,721	1,045,443	975,290	259,988
Transaction Size (\$)	354,067	304,030	552,236	210,209
Coupon (%)	5.50	5.42	5.58	5.58
Age	5.48	4.97	6.23	5.85
Maturity	14.42	15.42	13.2	13.35
Modified Duration	9.53	9.83	8.48	10.16

Table 2: Average number of trades per month for different levels of seasoning. Seasoning is computed as the time since issuance as a percentage of the original term of the bond. The second row in the table represents average number of trades per month conditional on the bond trading in that decile of its life.

Seasoning	Number of Transactions	
	Unconditional	Conditional
< 10%	5.87	15.75
10 – 20%	0.16	1.34
20 – 30%	0.13	1.04
30 – 40%	0.11	0.93
40 – 50%	0.09	0.81
50 – 60%	0.07	0.78
60 – 70%	0.06	0.77
70 – 80%	0.06	0.78
80 – 90%	0.05	0.75
≥ 90%	0.04	0.67

Table 3: Markups in basis points for different transaction sizes. The first line reports the average markup. The second line reports the standard deviation.

Par	All	Immediate Match	Not Overnight	Overnight	Not Split	Split	No Interdealer	Interdealer
All	208.9 158.2	159.1 136.6	165.1 142.6	227.8 160.9	209.0 160.6	208.8 149.0	200.5 146.7	218.2 175.8
(0, 0.005]	309.9 170.6	253.3 140.2	264.2 151.7	331.6 174.8	309.9 170.6	170.5 180.2	291.8 153.9	388.7 204.7
(0.005, 0.01]	262.6 146.2	215.7 121.1	223.0 128.8	279.8 149.9	261.5 146.0	293.8 147.1	247.6 129.6	307.5 179.7
(0.01, 0.025]	245.5 142.0	196.1 119.2	206.5 125.5	262.0 145.3	239.3 141.4	283.4 139.9	232.3 127.3	280.2 172.4
(0.025, 0.05]	210.9 138.0	164.1 117.3	177.3 124.7	223.5 140.5	197.4 135.6	257.6 135.9	201.4 126.9	219.0 152.0
(0.05, 0.1]	182.5 138.8	124.9 113.1	145.4 123.3	195.0 141.5	156.1 132.1	225.6 138.7	174.8 129.2	177.4 142.9
(0.1, 0.25]	140.8 136.1	72.0 93.2	93.7 108.7	157.6 140.9	105.0 120.2	188.3 141.4	128.8 128.4	135.6 129.3
(0.25, 0.5]	107.6 124.6	40.2 66.4	57.4 83.7	127.3 132.3	64.5 94.9	147.8 135.1	87.5 115.0	106.9 106.9
(0.5, 0.75]	83.8 115.5	24.5 52.1	37.8 68.2	108.5 127.5	40.6 77.8	131.3 130.6	55.0 94.4	99.6 113.1
(0.75, 1]	71.9 123.3	12.7 35.5	22.6 50.1	100.9 142.8	27.4 75.7	122.8 145.6	37.7 87.8	106.2 133.5
(1, 2.5]	42.0 87.4	9.6 29.8	15.1 40.6	66.5 108.7	16.4 52.3	98.5 117.3	19.2 59.1	71.9 96.2
(2.5, 5]	16.9 56.0	4.6 18.5	6.7 22.5	29.7 78.3	6.4 29.4	47.4 92.1	6.5 29.0	38.6 64.5
(5, ∞)	7.5 37.9	3.6 18.9	4.5 21.2	11.3 51.6	4.6 29.9	16.0 53.8	4.3 28.7	24.5 62.1

Table 4: Number of round-trip transactions, frequency of losses on round-trip transactions, median profit to dealers on a round-trip transaction, and median time in dealer inventory.

Par	Number of Trades	Frequency of Loss (%)	Dealer Profits (\$)	Time In Dealer Inventory (Days)
All	908,638	1.37	371	1.02
(0, 0.005]	113,948	0.67	148	0.86
(0.005, 0.01]	157,898	0.51	244	0.93
(0.01, 0.025]	131,413	0.43	392	0.95
(0.025, 0.05]	187,344	0.60	527	1.07
(0.05, 0.1]	119,095	0.84	848	1.41
(0.1, 0.25]	97,259	1.14	1,264	1.70
(0.25, 0.5]	26,323	1.36	2,076	2.02
(0.5, 0.75]	14,272	1.62	2,281	1.45
(0.75, 1]	3,603	1.05	1,527	1.23
(1, 2.5]	29,617	7.58	291	0.63
(2.5, 5]	12,173	10.86	0	0.21
(5, ∞)	15,693	18.31	0	0.22

Table 5: Summary statistics for the regressors. The first panel contains common-sample statistics for all purchase transactions in seasoned bond issues. The number of observations is adjusted for missing data. The second and third panels report common-sample statistics for the broadly and narrowly defined sample of matched pairs.

	Mean	Median	S. D.	Minimum	Maximum
All Transactions					
$\ln(Par)$	-2.97	-3.51	1.87	-6.91	5.99
Maturity	13.55	12.57	8.68	0.50	35.40
Coupon	4.34	5.13	2.46	0.00	16.00
Age	6.16	5.68	4.79	0.25	49.49
Transaction Frequency	0.15	0.05	0.32	0.00	4.84
Issue Size	-0.77	-1.20	2.22	-6.91	5.99
Volume	0.40	0.21	0.45	0.00	2.69
Order Imbalance	0.13	0.05	0.23	-0.73	2.69
New Issues	0.03	0.01	0.05	0.00	0.90
Treasury Yield	5.52	5.62	0.55	3.56	6.81
Treasury Range	6.48	5.58	3.88	0.00	39.53
Observations					1, 702, 285
Broad Sample					
$\ln(Par)$	-3.37	-3.69	1.62	-6.91	4.87
Maturity	11.46	9.97	8.00	0.50	35.40
Coupon	4.59	5.20	2.23	0.00	16.00
Age	6.35	6.15	4.86	0.25	49.49
Transaction Frequency	0.08	0.03	0.17	0.00	4.84
Issue Size	-1.54	-1.97	1.99	-6.91	5.99
Volume	0.39	0.20	0.44	0.00	2.69
Order Imbalance	0.13	0.05	0.23	-0.73	2.69
New Issues	0.03	0.01	0.05	0.00	0.90
Treasury Yield	5.49	5.60	0.59	3.56	6.81
Treasury Range	6.52	5.55	4.02	0.00	39.53
Observations					908, 638
Narrow Sample					
$\ln(Par)$	-3.32	-3.69	1.89	-6.91	4.64
Maturity	12.30	11.09	8.59	0.50	35.40
Coupon	4.49	5.25	2.42	0.00	15.50
Age	6.80	6.89	4.78	0.25	47.85
Transaction Frequency	0.12	0.04	0.26	0.00	4.84
Issue Size	-1.20	-1.61	2.23	-6.91	5.99
Volume	0.39	0.20	0.44	0.00	2.69
Order Imbalance	0.13	0.05	0.23	-0.73	2.69
New Issues	0.03	0.01	0.05	0.00	0.90
Treasury Yield	5.48	5.58	0.59	3.56	6.81
Treasury Range	6.45	5.50	4.00	0.00	39.53
Observations					221, 032

Table 6: Parameter Estimates of the Stochastic Frontier Model: Narrow Sample

	No Bargaining Power		Bargaining Power	
Cost Function Parameters				
Constant	170.13	(330.94)	2.08	(13.04)
$\ln(Par)$	-34.59	(-260.62)	0.65	(23.48)
Maturity	3.69	(130.82)	0.04	(9.41)
Coupon	-12.85	(-44.63)	-0.84	(-13.34)
Zero Coupon Bond	-84.81	(-49.02)	-8.76	(-22.76)
Age	-0.09	(-1.75)	0.03	(3.29)
Transaction Frequency	-14.34	(-22.39)	0.50	(7.82)
Issue Size	-3.81	(-26.16)	-1.11	(-30.19)
Volume	11.28	(19.26)	0.70	(8.88)
Order Imbalance	-1.43	(-1.45)	-0.25	(-1.79)
New Issues	-2.96	(-0.68)	-0.49	(-0.81)
Treasury Yield	10.28	(28.04)	0.19	(2.81)
Treasury Range	0.36	(7.06)	-0.02	(-2.24)
Revenue Bond	2.26	(4.74)	-0.08	(-1.04)
Tax Revenue Bond	2.86	(2.68)	0.10	(0.56)
General Obligation Bond	-3.43	(-2.42)	-0.87	(-4.44)
Guaranteed	22.28	(5.98)	-3.71	(-3.24)
Prerefunded	33.44	(23.46)	-0.48	(-1.52)
Certificates of Partnership	2.76	(2.76)	-0.08	(-0.54)
County Issuer	-0.41	(-0.86)	0.06	(0.86)
School District Issuer	1.04	(1.47)	0.14	(1.15)
Development Authority	-6.13	(-6.60)	-0.14	(-1.22)
Financial Authority	-5.06	(-4.61)	-0.26	(-1.85)
Housing Authority	-10.97	(-4.73)	-0.75	(-2.33)
Sewer Authority	5.66	(5.59)	0.14	(0.81)
Utilities	7.58	(6.23)	0.86	(4.44)
Redevelopment	14.33	(7.24)	-1.16	(-2.66)
Improvement	-2.55	(-2.99)	0.00	(0.01)
Facilities	-0.38	(-0.60)	0.32	(3.59)
Refinancing	6.44	(14.12)	0.41	(5.89)
Symmetric Error Distribution				
Constant	101.50	(3070.95)	9.92	(180.83)
$\ln(Par)$	-0.14	(-178.47)	-0.31	(-106.00)
Asymmetric Error Distribution				
Constant	-	-	123.35	(2133.11)
$\ln(Par)$	-	-	-0.52	(-380.56)
Log Likelihood	-1,333,000		-1,296,827	

Table notes on the next page.

Notes for Table 6

Table 6 reports maximum likelihood estimates from fitting the model

$$p_{ij} - p_{ij}^* = \theta_o + \sum_{l=1}^L \theta_l X_l + \varepsilon_i + u_{ij},$$

with p_{ij} the price that the dealer sells the bond for, p_{ij}^* the price that the dealer pays the seller for the bond, X_l for $l = 1, \dots, L$ conditioning variables. The residual ε_i is normally distributed with standard deviation $b_0 \prod_{k=1}^K e^{b_k Z_k}$, and u_{ij} is exponential distributed with parameters $a_0 \prod_{k=1}^K e^{a_k Z_k}$, with Z_k for $k = 1, \dots, K$ conditioning variables. The first two columns report the parameter estimates and t-statistics in parentheses for the model with $a_0 = a_1 = \dots = a_K = 0$ and the second two columns report the parameter estimates and t-statistics in parentheses for the unrestricted model. The number of observations is 221,032.

Table 7: Parameter Estimates of the Stochastic Frontier Model: Broad Sample

	No Bargaining Power		Bargaining Power	
Cost Function Parameters				
Constant	153.11	(409.68)	2.85	(17.02)
$\ln(Par) \times Par \in (0, 0.1)$	-55.86	(-281.59)	-77.24	(-298.50)
$\ln(Par) \times Par \in [0.1, 0.5)$	-51.97	(-142.23)	-2.58	(-21.58)
$\ln(Par) \times Par \in [0.5, \infty)$	-39.17	(-308.89)	0.73	(25.16)
Overnight $\times Par \in (0, 0.1)$	45.57	(145.04)	64.79	(196.53)
Overnight $\times Par \in [0.1, 0.5)$	35.13	(50.49)	6.09	(24.47)
Overnight $\times Par \in [0.5, \infty)$	1.36	(2.02)	-1.03	(-14.42)
Days In Inventory \times Overnight $\times Par \in (0, 0.1)$	-0.41	(-13.73)	-2.09	(-67.35)
Days In Inventory \times Overnight $\times Par \in [0.1, 0.5)$	-0.84	(-14.13)	-0.85	(-33.57)
Days In Inventory \times Overnight $\times Par \in [0.5, \infty)$	0.33	(4.10)	-0.43	(-28.94)
Split \times No Interdealer $\times Par \in (0, 0.1)$	44.10	(103.31)	83.25	(117.86)
Split \times No Interdealer $\times Par \in [0.1, 0.5)$	72.67	(90.26)	11.91	(21.40)
Split \times No Interdealer $\times Par \in [0.5, \infty)$	36.48	(42.80)	-0.95	(-11.19)
No Split \times Interdealer $\times Par \in (0, 0.1)$	65.85	(119.51)	27.56	(47.96)
No Split \times Interdealer $\times Par \in [0.1, 0.5)$	43.68	(46.79)	11.74	(26.11)
No Split \times Interdealer $\times Par \in [0.5, \infty)$	56.69	(45.99)	-2.28	(-6.16)
Split \times Interdealer $\times Par \in (0, 0.1)$	84.36	(83.29)	69.72	(57.79)
Split \times Interdealer $\times Par \in [0.1, 0.5)$	89.87	(88.90)	26.99	(45.14)
Split \times Interdealer $\times Par \in [0.5, \infty)$	111.44	(105.13)	3.92	(8.09)
Maturity	6.00	(337.07)	0.15	(34.37)
Coupon	-7.43	(-45.27)	-1.59	(-26.39)
Zero Coupon Bond	-18.04	(-18.51)	-14.42	(-39.30)
Age	-0.49	(-16.60)	-0.03	(-3.89)
Transaction Frequency	-40.73	(-65.89)	0.03	(0.42)
Issue Size	-4.09	(-48.62)	-1.06	(-32.78)
Volume	16.26	(46.17)	0.71	(12.17)
Order Imbalance	0.87	(1.50)	0.07	(0.65)
New Issues	5.93	(2.26)	1.03	(2.38)
Treasury Yield	11.22	(53.46)	1.62	(25.31)
Treasury Range	0.58	(19.50)	0.05	(7.72)
Revenue Bond	1.83	(6.52)	0.14	(2.21)
Tax Revenue Bond	3.56	(5.76)	0.30	(2.10)
General Obligation Bond	1.70	(1.89)	-0.82	(-4.93)
Guaranteed	13.47	(6.62)	0.24	(0.36)
Prerefunded	42.86	(49.49)	3.83	(9.16)
Certificates of Partnership	1.81	(3.07)	0.27	(2.27)
County Issuer	-0.81	(-2.87)	-0.02	(-0.37)
School District Issuer	2.87	(7.32)	0.46	(4.38)
Development Authority	-5.73	(-9.32)	-0.24	(-2.58)
Financial Authority	-7.33	(-10.23)	-0.56	(-5.07)
Housing Authority	-6.15	(-4.57)	-0.64	(-2.85)
Sewer Authority	2.98	(5.03)	0.51	(3.32)
Utilities	4.81	(6.79)	0.74	(4.37)
Redevelopment	11.58	(10.76)	1.18	(4.04)
Improvement	-3.09	(-6.06)	-0.23	(-2.23)
Facilities	-0.08	(-0.20)	0.37	(5.30)
Refinancing	4.33	(16.08)	0.35	(6.18)

The table is continued on the next page.

Table 7: Parameter Estimates of the Stochastic Frontier Model: Broad Sample (continued)

	No Bargaining Power		Bargaining Power	
Symmetric Error Distribution				
Constant	100.93	(2578.05)	30.76	(616.34)
$\ln(Par) \times Par \in (0, 0.1)$	-0.13	(-111.12)	-0.57	(-181.98)
$\ln(Par) \times Par \in [0.1, 0.5)$	-0.05	(-19.15)	-0.88	(-89.23)
$\ln(Par) \times Par \in [0.5, \infty)$	-0.09	(-94.94)	-0.51	(-249.11)
Overnight $\times Par \in (0, 0.1)$	0.02	(13.01)	0.24	(49.10)
Overnight $\times Par \in [0.1, 0.5)$	0.12	(25.92)	-0.05	(-2.29)
Overnight $\times Par \in [0.5, \infty)$	0.08	(14.51)	0.34	(27.51)
Days In Inventory \times Overnight $\times Par \in (0, 0.1)$	0.01	(82.30)	0.00	(5.11)
Days In Inventory \times Overnight $\times Par \in [0.1, 0.5)$	0.01	(40.68)	0.01	(10.00)
Days In Inventory \times Overnight $\times Par \in [0.5, \infty)$	0.02	(45.03)	0.02	(20.22)
Split \times No Interdealer $\times Par \in (0, 0.1)$	-0.01	(-1.95)	0.58	(76.45)
Split \times No Interdealer $\times Par \in [0.1, 0.5)$	0.01	(1.53)	0.83	(22.71)
Split \times No Interdealer $\times Par \in [0.5, \infty)$	0.16	(21.65)	0.00	(0.12)
No Split \times Interdealer $\times Par \in (0, 0.1)$	0.37	(149.25)	0.25	(34.65)
No Split \times Interdealer $\times Par \in [0.1, 0.5)$	0.04	(5.89)	0.65	(25.01)
No Split \times Interdealer $\times Par \in [0.5, \infty)$	0.04	(4.28)	0.58	(19.74)
Split \times Interdealer $\times Par \in (0, 0.1)$	0.35	(74.02)	0.53	(32.74)
Split \times Interdealer $\times Par \in [0.1, 0.5)$	0.10	(16.88)	1.17	(44.00)
Split \times Interdealer $\times Par \in [0.5, \infty)$	0.12	(14.69)	1.01	(39.47)
Asymmetric Error Distribution				
Constant	-	-	141.49	(1710.83)
$\ln(Par) \times Par \in (0, 0.1)$	-	-	0.12	(44.59)
$\ln(Par) \times Par \in [0.1, 0.5)$	-	-	-0.44	(-108.90)
$\ln(Par) \times Par \in [0.5, \infty)$	-	-	-0.62	(-389.74)
Overnight $\times Par \in (0, 0.1)$	-	-	-0.14	(-36.63)
Overnight $\times Par \in [0.1, 0.5)$	-	-	0.26	(36.73)
Overnight $\times Par \in [0.5, \infty)$	-	-	0.33	(34.74)
Days In Inventory \times Overnight $\times Par \in (0, 0.1)$	-	-	0.01	(40.46)
Days In Inventory \times Overnight $\times Par \in [0.1, 0.5)$	-	-	0.00	(3.37)
Days In Inventory \times Overnight $\times Par \in [0.5, \infty)$	-	-	0.02	(38.71)
Split \times No Interdealer $\times Par \in (0, 0.1)$	-	-	-0.16	(-23.23)
Split \times No Interdealer $\times Par \in [0.1, 0.5)$	-	-	0.57	(67.12)
Split \times No Interdealer $\times Par \in [0.5, \infty)$	-	-	0.92	(82.08)
No Split \times Interdealer $\times Par \in (0, 0.1)$	-	-	0.29	(53.18)
No Split \times Interdealer $\times Par \in [0.1, 0.5)$	-	-	0.22	(22.99)
No Split \times Interdealer $\times Par \in [0.5, \infty)$	-	-	1.11	(67.81)
Split \times Interdealer $\times Par \in (0, 0.1)$	-	-	0.21	(21.03)
Split \times Interdealer $\times Par \in [0.1, 0.5)$	-	-	0.53	(55.60)
Split \times Interdealer $\times Par \in [0.5, \infty)$	-	-	1.56	(121.40)
Log Likelihood	-5,619,000		-5,509,721	

Table notes on the next page.

Notes for Table 7

Table 7 reports maximum likelihood estimates from fitting the model

$$p_{ij} - p_{ij}^* = \theta_o + \sum_{l=1}^L \theta_l X_l + \varepsilon_i + u_{ij},$$

with p_{ij} the price that the dealer sells the bond for, p_{ij}^* the price that the dealer pays the seller for the bond, X_l for $l = 1, \dots, L$ conditioning variables. The residual ε_i is normally distributed with standard deviation $b_0 \prod_{k=1}^K e^{b_k Z_k}$, and u_{ij} is exponential distributed with parameters $a_0 \prod_{k=1}^K e^{a_k Z_k}$, with Z_k for $k = 1, \dots, K$ conditioning variables. The first two columns report the parameter estimates and t-statistics in parentheses for the model with $a_0 = a_1 = \dots = a_K = 0$ and the second two columns report the parameter estimates and t-statistics in parentheses for the unrestricted model. The number of observations is 908,638.

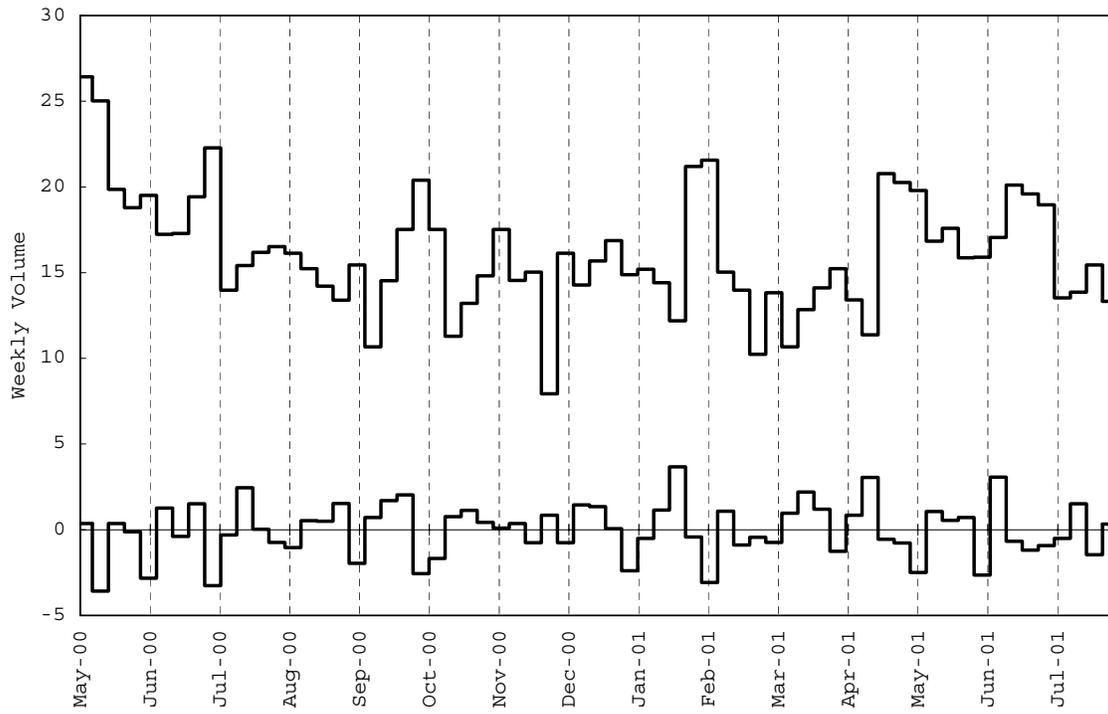


Figure 1: Volume of Trade and Order Imbalance in Seasoned Issues. The top line is sales to customers of seasoned issues and the bottom line is the aggregate order imbalance—the difference between sales to and purchases from customers.

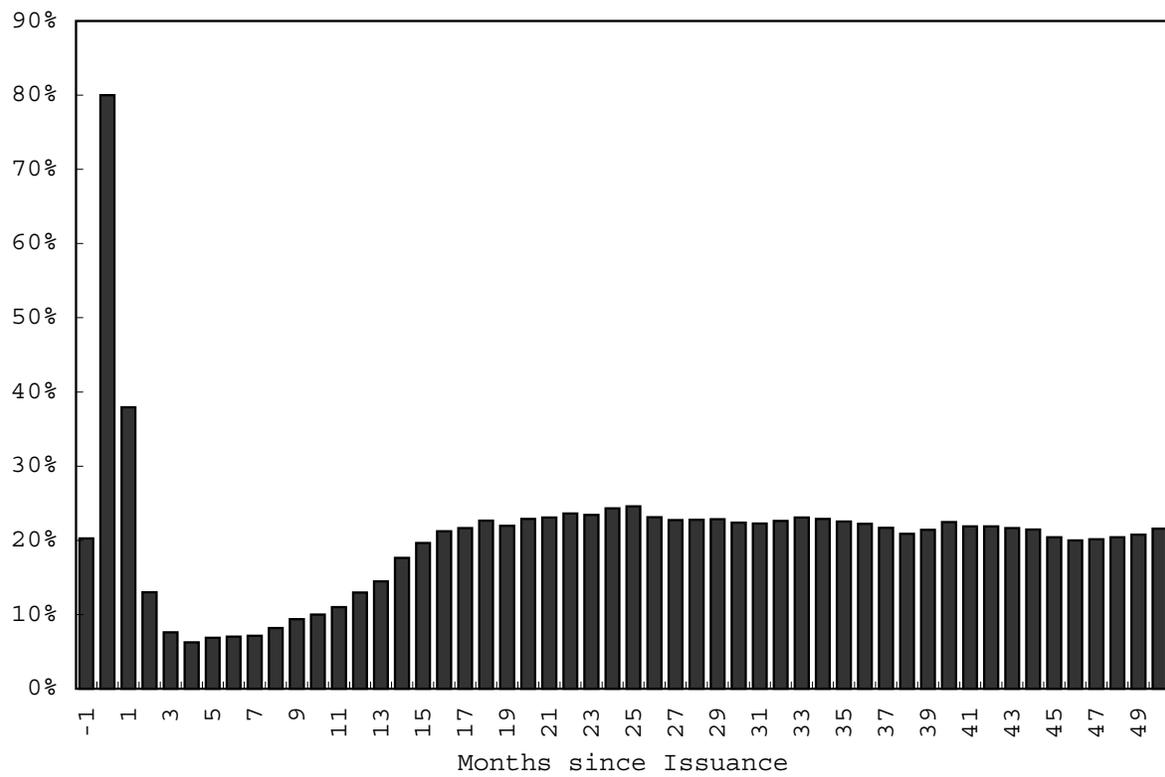
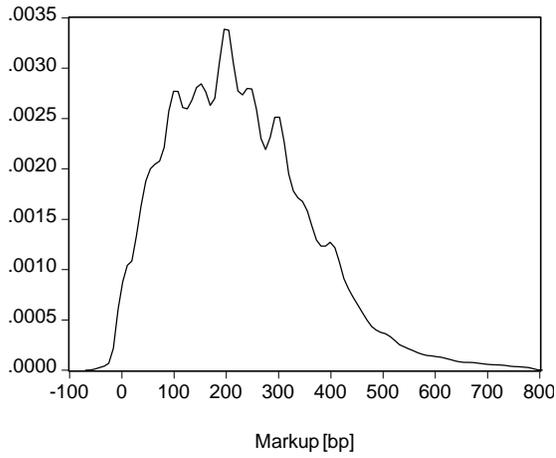


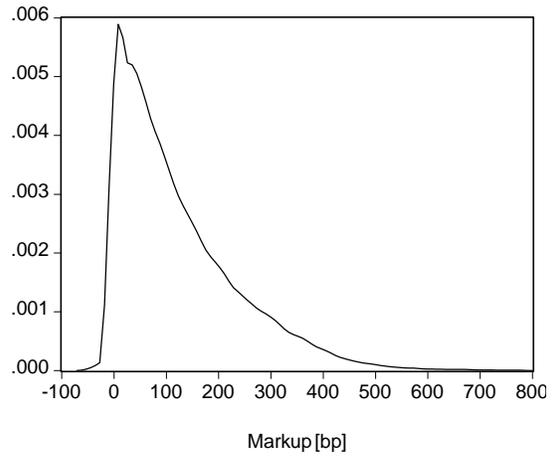
Figure 2: Average probability of trade in a month



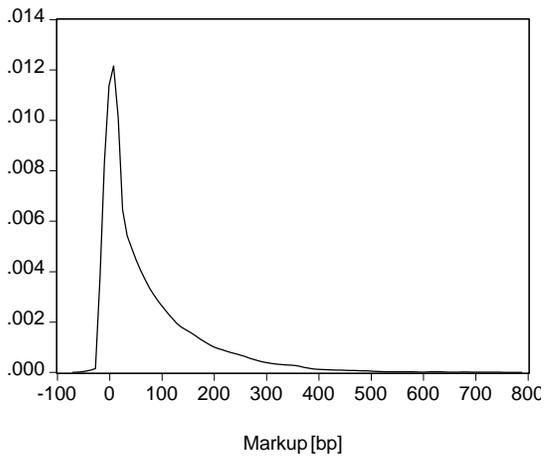
Figure 3: Average number of trades in a seasoned bond per day conditional on the bond trading on that day.



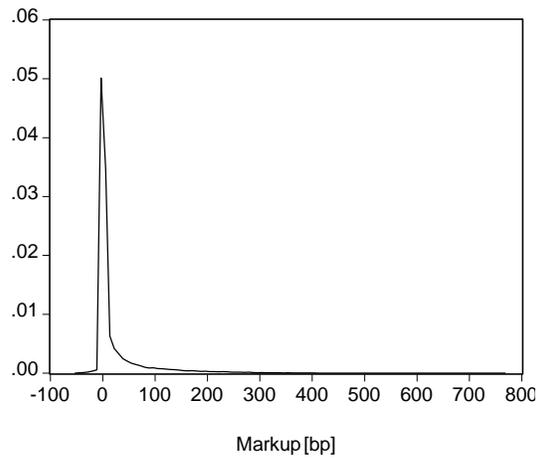
(a) Panel A: $Par < 100,000$.



(b) Panel B: $100,000 \leq Par < 500,000$.



(c) Panel C: $500,000 \leq Par < 1,000,000$.



(d) Panel D: $Par \geq 1,000,000$.

Figure 4: Distribution of markups for different transaction sizes. The distributions are estimated using Epanechnikov kernels with the Silverman automatic bandwidth selection.