

Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall?

RICHARD C. GREEN, DAN LI, and NORMAN SCHÜRHOFF*

ABSTRACT

We study price discovery in municipal bonds, an important OTC market. As in markets for consumer goods, prices “rise faster than they fall.” Round-trip profits to dealers on retail trades increase in rising markets but do not decrease in falling markets. Further, effective half-spreads increase or decrease more when movements in fundamentals favor dealers. Yield spreads relative to Treasuries also adjust with asymmetric speed in rising and falling markets. Finally, intraday price dispersion is asymmetric in rising and falling markets, as consumer search theory would predict.

IN A WIDE VARIETY OF CONSUMER MARKETS, prices for retail goods rise faster than they fall.¹ This asymmetric price adjustment, referred to as “rockets and feathers,” is generally understood by economists to be inconsistent with perfect competition between sellers. Sellers appear to exploit local market power due to the search frictions that their customers face.

We document asymmetric price adjustment in a major over-the-counter (OTC) financial market using a comprehensive sample of all trades in municipal bonds over a 5-year period. We provide evidence that the asymmetric price adjustment is associated with opportunistic timing by the broker-dealers who intermediate trades in the market. Movements in fundamentals that increase spreads earned by broker-dealers are reflected quickly in prices, while prices respond more slowly to changes that reduce their spreads. Further, the asymmetries are more pronounced in states where tax treatment of the bonds would induce market segmentation. We also document asymmetries in within-day price dispersion, as implied by theoretical models of rockets-and-feathers price adjustments based on costly consumer search and learning.

*Richard C. Green is with the Tepper School of Business at Carnegie Mellon University, Dan Li is with the Board of Governors of the Federal Reserve System, and Norman Schürhoff is with the Faculty of Business and Economics at University of Lausanne, Swiss Finance Institute, and CEPR. The third author acknowledges research support from the Swiss Finance Institute and from NCCR FINRISK of the Swiss National Science Foundation. Seminar participants at Arizona State, Carnegie Mellon, Columbia, George Mason, Lausanne, Miami, Oklahoma, Oxford, and Wharton provided helpful comments and guidance. We greatly appreciate the feedback from our discussants Larry Harris, Adam Kolasinski, and Bruce Lehmann and conference participants at the AFA 2009, EFA 2008, and 2nd Workshop on Financial Market Quality. An anonymous referee and Associate Editor, along with the Editor, Campbell Harvey, have done much to improve the paper. We thank Tal Heppenstall, J.C. Stille, and Jim Konieczny of University of Pittsburgh Medical Center for conversations that alerted us to the problems studied in this paper.

¹See Peltzman (2000) and the many references therein.

In the municipal bond market, unlike the markets for most consumer goods, dealers trade with retail customers as both buyers and sellers. We show that on the ask side of the market, where dealers are selling, prices rise faster than they fall. On the bid side, where dealers are buying, prices fall faster than they rise. Since dealers in seasoned municipals tend to intermediate between large, better informed sellers, on the one hand, and small investors who buy and hold, on the other, measures of average prices such as the midpoint inherit the asymmetry on the ask side; that is, they rise faster than they fall, just as prices in consumer markets do. On average dealers are “buying wholesale” and “selling retail,” so the asymmetric movement in prices benefits them.

Our findings are of obvious importance to investors, issuers, and regulators involved with the municipal bond market, which represents \$1.7 trillion of outstanding bonds. The implications of these findings reach beyond the immediate setting, however, for at least three reasons.

First, many securities trade in opaque, decentralized OTC markets. Regulators and investors often complain that the lack of price transparency in these markets interferes with price discovery, discourages trade, increases costs of trade, and creates local market power for financial intermediaries. The direction of causality, however, is ambiguous as trading venue, volume of trade, and costs of trade are jointly endogenous outcomes. For instance, the securities in question may trade in decentralized markets precisely because investors rarely wish to trade them. Dealers might therefore earn high profits arranging trades, even under perfect competition, if the costs of identifying and matching counterparties with coincident needs are high. Our findings of asymmetric price adjustment are instructive because it is difficult to imagine any more primitive condition that would lead prices to rise faster than they fall in a competitive marketplace. The only explanations of such behavior offered in the economics literature involve search costs for consumers and imperfect competition between sellers.

Second, our research suggests that the market is failing at one of its fundamental functions, efficient price discovery. Past research on OTC financial markets shows that the costs of trade are high for retail investors, and that these higher costs appear to be due to search costs and a lack of price transparency.² An unanswered question is how these costs affect price discovery—the efficiency with which prices reflect information. Even if trade is costly and infrequent, the transactions that do occur can still be at prices that efficiently incorporate public information. If retail investors face high costs of trade, financial intermediaries such as mutual funds will arise to limit these costs. Thus, investors who are unsophisticated shoppers for municipal or corporate bonds are only hurting themselves. Financial market prices, however, have consequences that reach beyond the investors and intermediaries directly involved in a given trade. Prices serve as signals for resource allocation more

²See, for example, Harris and Piwowar (2006), Bessembinder, Maxwell, and Venkataraman (2006), Goldstein, Hotchkiss, and Sirri (2007), Edwards, Harris, and Piwowar (2007), and Green, Hollifield, and Schürhoff (2007a).

broadly. Our results show that transaction prices in the municipal market respond very slowly to information that moves interest rates generally. Indeed, at times prices fail to move at all for days after an important announcement, and when prices do adjust, they do so asymmetrically in rising and falling markets.

Third, our paper is the first to show that a behavior known to be pervasive in consumer markets also characterizes financial asset prices.³ This is instructive about the magnitude and nature of the search frictions that investors face. Consumer markets are often characterized by high levels of product differentiation, and stores are geographically disperse. Further, many consumer goods have private-value attributes. Search costs are therefore heterogeneous across consumers for reasons that are easy to understand. Financial securities, in contrast, are common-value assets—that is, agents with the same information ascribe the same intrinsic value to them. Further, no transportation or physical examination is required to locate and evaluate a municipal bond, so search costs are entirely informational in nature. Thus, if we observe asymmetric price adjustment for financial assets, we can more confidently attribute it to a lack of price transparency. The differential search costs are, in turn, easier to alter or eliminate through regulatory choices. Our results on the municipal bond market show that these costs alter the dynamic evolution of prices, and thus efforts to improve transparency and disclosure can improve price discovery.

Most studies of the rockets-and-feathers phenomenon are based on data that are of high frequency on the time-series dimension. The municipal market, in contrast, involves a huge cross section (over 1.5 million separate bonds), but trading in individual bonds is relatively infrequent. We therefore propose statistical models that aggregate all transactions in a bond at the daily level, and we construct proxies for effective bid-ask spreads, half-spreads,⁴ and yield spreads over comparable Treasuries. For a given bond on a given day, we estimate the value of the bond as the midpoint between the lowest price (highest yield) at which a dealer sold to a customer and the highest price (lowest yield) at which a dealer bought from a customer. If we do not observe both sales and buys on a given day, we use the average price for all interdealer transactions. The sale price on that day is then the average price at which bonds were sold to customers, and analogously for the buy price. We then use panel data methods to examine price discovery and implicit bid-ask spreads.

Our panel regressions show that, while Treasury rates respond quickly to macroeconomic movements, municipal rates do not. Yield spreads also respond dramatically and persistently because the price adjustment for municipals is so slow. We find that municipal yields respond sluggishly to news in part because they often do not respond at all for several days. We find no evidence that volume is significantly higher on news days than on other days.

We further show that prices adjust asymmetrically in a number of ways. First, we limit attention to dealer purchases we can match with subsequent

³Some evidence for retail banking is documented in Neumark and Sharpe (1992).

⁴Half-spreads are the difference between the sales price and the midpoint or the purchase price and the midpoint, normalized by the midpoint.

sales. We find that profits on retail trades rise in rallies faster than they decrease when prices fall. This evidence suggests that dealers opportunistically delay recognition of movements in fundamentals.

Next, we ask how effective half-spreads respond to changes in the midpoint, our proxy for the bond's value. When the midpoint of the bond rises, the average sales price less the midpoint (the ask-side effective spread) is unaffected, but it rises when values fall. Thus, on the ask side, prices rise faster than they fall. Similarly, the bid-side effective half-spread rises when prices rise, but is unaffected when prices fall. Thus, the prices that dealers pay fall faster than they rise. In short, when underlying values move to dealers' advantage, dealers quickly adjust prices up or down to maintain a constant profit margin. In contrast, when price movements decrease the cost of a bond that dealers are selling, or increase the cost of a bond they are buying, dealers' transactions prices are sticky.

Evidence of asymmetric price responses is also seen in the speed with which bid-ask spreads and half-spreads adjust to changing conditions over time. Using a partial-adjustment model with random effects for bid-ask spreads and half-spreads, we find that when effective spreads are narrow (relative to the latent "equilibrium" value) they rise immediately, whereas when these spreads are wide, the adjustment is more gradual.

Next, we consider the movements over time of yield spreads between municipals and Treasuries using a partial adjustment model. For each municipal bond for which we have estimates of its value on consecutive days, we consider the yield spread relative to a maturity-matched Treasury. When the spread of the Treasury yield over the municipal's midpoint yield is abnormally high relative to a latent "equilibrium" value, the municipal's yield tends to rise, and thus the price tends to fall. The reverse occurs when the yield spread is unusually narrow. Our estimates show that yield spreads widen faster than they shrink. Alternatively stated, municipal prices (ask and midpoint prices) rise faster than they fall, as we would expect if the composition of customers is asymmetric across the two sides of the market. Indeed, on average dealers are buying from larger, better-informed investors than they are selling to. Thus, the asymmetries associated with the ask side dominate the movement of "the price."

In our final tests, we use variation across states to verify that asymmetric price adjustment is more pronounced in settings where we would a priori expect dealers to have more local monopoly power, and we test previously unstudied implications of theoretical models of asymmetric price adjustment. This analysis shows that within-day dispersion in the prices at which dealers sell to customers is higher when prices are falling than when they are rising, as the theory predicts. We find the exact opposite for prices at which dealers buy from customers. Taken together, these findings suggest that dealers are exploiting search frictions on both sides of the market.

The remainder of the paper is organized as follows. Section I describes the data used. Section II discusses the institutional setting in more detail. In Section III we evaluate the speed with which the municipal market responds to

macroeconomic news, and show that effective bid-ask spreads and prices adjust asymmetrically to movements in fundamentals. In Section IV we consider how these behaviors vary across states and we examine intraday price dispersion. Section V summarizes and concludes.

I. Data

We study price discovery in the municipal bond market using data provided by the Municipal Securities Rulemaking Board (MSRB), a self-regulatory industry group. These data include all trades made by registered broker-dealers in municipal securities from May 1, 2000 to October 19, 2006. This is essentially the entire secondary market for municipal bonds.⁵ There are 1,615 trading days during this period. Trades are reported in 1,559,894 bonds. Dealers are required to record, for all trades, the time, the par value, and whether the trade was a purchase from a customer, a sale to a customer, or an interdealer trade. Individual dealers are not identified in the data.

We apply a number of rule-based filters to clean the transactions data, eliminate bonds with missing observations, correct obvious clerical errors, and supply missing data items where possible. We exclude a small number of trades on holidays and weekends. For each bond issue, we search the MSRB database for the bond's coupon rate and maturity, since these are reported for some trades but missing for others. We eliminate all bonds with missing coupon or maturity after this search. This includes all variable-rate bonds. Bonds with a recorded coupon in excess of 20% or a maturity of more than 100 years are dropped as likely to have been misrecorded. Our panel regressions require a minimal amount of time-series information, and so for these purposes we eliminate any bonds with fewer than 10 trades over the sample period. Our tests involving round-trip transactions do not impose this filter, as they require much more limited time-series information.

Green, Hollifield, and Schürhoff (2007b) show that newly issued bonds exhibit some peculiar behaviors and high levels of price dispersion. We therefore focus on seasoned bond issues in this study and eliminate transactions less than 180 days after the bonds are issued.⁶ As shown in Green, Hollifield, and Schürhoff (2007a), by this point dealer purchases from and sales to customers are roughly equal in value. Our sample after applying the various filters has sales to customers of \$1.42 trillion, purchases from customers of \$1.37 trillion, and interdealer trades of \$1.17 trillion in par value. The number of sales to customers, however, is much larger than dealer purchases. Dealers often provide liquidity to institutional traders by purchasing large blocks of bonds and selling smaller blocks off to retail investors, or to regional dealers with retail

⁵There are no mechanisms in place through which investors can trade directly with each other. Even newer electronic platforms such as MuniCenter, which allow institutions to directly negotiate a price, treat the buyer and seller as anonymous, with the platform (a broker-dealer) acting as counterparty to each.

⁶Specifically, 180 days after the "dated date," when interest begins to accrue.

Table I
Macroeconomic News Events

Macro news is measured by the standardized surprise component in the macro announcement. We compute the standardized surprise, as in Balduzzi, Elton, and Green (2001), as the actual value minus the consensus forecast divided by the standard deviation of surprises across time. A “–” in the fourth column indicates that bond prices fall on average in response to a positive surprise, and a “+” indicates that they rise.

	Events	Announcement Time	Effect on Bond Prices	Average Positive Surprise	Average Negative Surprise
Advance retail sales	81	8:30 am	–	0.79	–0.61
Capacity utilization	80	9:15 am	–	0.92	–0.91
Change in nonfarm payrolls	81	8:30 am	–	0.59	–0.92
Consumer price index	33	8:30 am	–	1.20	–1.06
GDP annualized	78	8:30 am	–	0.78	–0.81
Industrial production	81	9:15 am	–	0.93	–0.93
Producer price index	47	8:30 am	–	0.98	–0.77
U. of Michigan confidence	157	9:45–10:00 am	–	0.65	–0.86
Unemployment rate	82	8:30 am	+	1.18	–1.12
FOMC rate decision	55	~2:15 pm	–	4.29	–4.29

distribution capability. Our sample has 12.90 million sales to customers, 7.04 million purchases from customers, and 6.22 million interdealer trades, for a total of 26.16 million trades.

Along with the information in the MSRB transactions data, we use data on issuer and bond characteristics from other sources. We employ the Thomson-Reuters SDC Platinum database on municipal securities, which covers new municipal issues from 1980 to 2006. The Financial Industry Regulatory Authority also provides data on bonds and issuers through the Securities Industry and Financial Markets Association’s web page, www.investingbonds.com. All explanatory variables that exhibit outliers in either tail are winsorized at 0.5% and 99.5% if they attain both positive and negative values, or at 99% if they are only defined for positive values.

To evaluate the price response to major news events, we employ a number of periodic macroeconomic news releases common in the literature. These are detailed in Table I and cover the period May 2000 through October 2006. In the table we also list whether a positive surprise in the variable in question is, on average, good news (+ sign) or bad news (– sign) for bond prices. We categorize these events based on intraday studies of the effects of announcements on Treasury bond prices, such as Balduzzi, Elton, and Green (2001).⁷ We make use of this information to sort announcements as good or bad news for bonds. For example, a negative surprise in advance retail sales is good news, and a positive surprise is bad news, while a positive surprise in the unemployment rate is bad news and a negative surprise is good news.

⁷See Table II on p. 530 of Balduzzi, Elton, and Green (2001).

When we compute yield spreads over Treasuries for specific bonds, we use the daily constant maturity Treasury rates provided by the St. Louis Federal Reserve. These rates are for maturities of 0.25, 1, 2, 3, 5, 7, 10, 20, and 30 years. For intermediate maturities we interpolate linearly.

There are a large number of bonds outstanding, but most individual bonds trade infrequently. Because of differences in the terms of trade across types of investors, intraday price variation can be large compared to movements in fundamentals (see, e.g., Green, Hollifield, and Schürhoff (2007a)). For many of our tests, therefore, we employ panel data methods and focus on transactions data aggregated at a daily frequency. Daily municipal yields for specific bonds are constructed by taking the midpoint of the highest yield on all sales by dealers to customers and the lowest yield on all purchases by dealers from customers of a particular bond on a given day. If there are not both sales and purchases on that day, the daily yield is the average (or median) yield on interdealer trades. We compute midpoints for prices similarly as the midpoint of the lowest offer and the highest bid price. In many of our tests we also compute changes in yields, or require lagged yields over several days. We have fewer observations for these tests since many bonds do not trade on consecutive days. Yield spreads are calculated by matching the yield on a given municipal bond with a Treasury yield of comparable maturity. Before running our tests using variables that rely on midpoints constructed from the inside spread, we truncate the distributions of those variables at the lowest and highest 0.5%.⁸

In the following sections we show that individual municipal bond prices adjust to new information very slowly, that they rise faster than they fall, and that these patterns are consistent with local market power for dealers and search frictions for investors. One would expect that opportunities to short sell would speed price adjustment when prices are falling. However, there are no mechanisms for shorting individual municipal bonds available to investors. There are CDS contracts for states and other very large issuers. Yet there is no mechanism through which a speculator can directly benefit from the price sluggishness we document for individual bonds without trading in that bond, or without circumventing the intermediaries and gaining access to the customers who are buying or selling at stale prices.

II. Dealer Intermediation in the Municipal Bond Market

Dealers in the market for seasoned municipal bonds provide liquidity for their customers either by taking bonds into inventory and searching for a buyer, or by arranging trades beforehand. In the former case, we would see an initial purchase from a customer in our data followed by subsequent sales to customers, potentially spread out over time. In the latter case we would see a

⁸In estimating the value of the bond, we encountered occasional situations in which the highest price paid by a dealer exceeds the lowest price the dealer received for the same bond. These outliers are filtered out by truncating the distribution of the inside spread at the lowest and highest 0.5%.

purchase and sale of the same bond and the same par value within minutes of each other.

Dealers do not provide liquidity through short positions. The infrequent trade in individual bonds makes a dealer's ability to cover a short position highly uncertain. Indeed, common usage of the term "short" in the municipal market retains its original meaning of "falling short" or "being in deficit." It refers to situations where, through clerical error or miscommunication, a dealer inadvertently sells a bond not in inventory. These situations are typically resolved by providing the customer with a similar bond and some "going-away money."⁹

While dealers face retail customers on both sides of the market, one of the major ways in which they provide liquidity to institutions such as hedge funds and mutual funds is by purchasing large blocks of bonds, and then selling these off in smaller pieces to retail investors or to regional broker-dealers with retail distribution capacity. These bonds would typically have been purchased by the institution when the bond was issued. As a result, the market for seasoned municipals has many of the attributes of a retail market, such as the gasoline market, where the intermediaries buy at wholesale prices and sell in smaller quantities to less sophisticated retail customers. This is evident in the patterns of trade in the data, where the par value of sales to customers, \$1.42 trillion, is roughly equal to the par value of purchases from customers, \$1.37 trillion, while the number of sales to customers, 12.90 million, is almost twice the number of purchases from them, 7.04 million. Thus, on average, purchases from customers are being broken in half as the dealers intermediate between institutional sellers and retail buyers. This pattern is also apparent from the distribution of transaction sizes. While 52% of all transactions are retail-sized sales to customers, dealer purchases from customers of up to \$50 thousand in par value constitute only 26% of all dealer purchases.

In Part A of the Internet Appendix to this paper we provide an example of the raw data from the MSRB database.¹⁰ It is evident from casual inspection of these data that purchases from customers generally precede sales to customers. We also provide additional evidence on the amount of time between trades that supports the view that dealers intermediate between investors by "buying wholesale and selling retail."

Suppose dealers can opportunistically delay the recognition of price movements in dealing with retail customers, while the prices at which they trade with institutions are more competitive. Then, if institutions are more active on the bid side of the market (selling to dealers), we will see asymmetries in prices and effective bid-ask spreads of the sort we document in our empirical analysis. The basic mechanism is as follows. The midpoint is the average of the highest price at which dealers purchase from customers and the lowest price at which dealers sell to customers on a given day. If institutions are actively trading a bond on that day, it is likely that the inside bid, and hence the midpoint of the inside spread, reflects their sales while the inside ask reflects retail trades. The

⁹We thank Chris Fama for bringing this to our attention.

¹⁰An Internet Appendix for this article is available online in the "Supplements and Datasets" section at <http://www.afajof.org/supplements.asp>.

Table II

Example of Asymmetric Response in Bid-Ask Spreads and Prices

The table reports quantities associated with our empirical analysis for a numerical example where intrinsic values rise from 100 to 104, or fall from 100 to 96.

	Prices Rise	Prices Fall
Initial value date 0	100	100
Intrinsic value date 1	104	96
Retail buyers' price	104	100
Institution seller's price	104	96
Retail seller's price	100	96
Midpoint v_{it}	104	98
Ave. cust. purchase price (ask)	104	100
Ave. cust. sale price (bid)	101.33	96
Ask-side half-spread (% of v_{it})	0.00	2.04
Bid-side half-spread (% of v_{it})	2.56	2.04

bid-side half-spread, in turn, is the midpoint less the average price at which dealers buy, while the ask-side half-spread is the average price at which dealers sell. If dealers can delay their response to movements in fundamentals when dealing with smaller traders, we would expect the ask-side spread to rise when prices fall, and the bid-side spread to rise when prices rise. On the other hand, when prices are falling dealers would quickly adjust the prices they pay downward, maintaining their spread; and, similarly, when prices are rising dealers would quickly adjust the prices at which they sell upward. This is exactly the pattern we document in Section III. Our central empirical findings are that prices, measured as daily midpoints of the inside implicit spread, rise faster than they fall (Table VIII), while half-spreads rise asymmetrically with the midpoint (Table VI).

A numerical example illustrates how our findings for implicit bid-ask spreads and for prices (or yields) can arise simultaneously, as long as institutions are more likely to be selling to dealers than buying from them, as is evident in our data. Assume that each period the dealer faces three customers with unit demand who buy, and three who sell. One of the customers who sells to the dealer is an institution, and the other two are retail investors. The institution always sells at the current intrinsic value, but the dealers can buy from the retail investors at the minimum of today's price and yesterday's price. The customers that buy from the dealer are all retail customers, and the dealer can charge these customers the maximum of today's price or yesterday's price. Now, suppose the intrinsic value of the security can either rise from 100 to 104, or fall from 100 to 96. If the price rises, all customers will buy at 104, which will be our measured ask price. The institution will demand the intrinsic value of 104 when he sells, but the retail sellers will only receive 100. If the price falls, the retail buyers will still pay 100, and both retail and institutional sellers will receive 96. Our measure of the value v_{it} is the midpoint of the inside spread. It and all the other quantities relevant to our results are computed in Table II.

The table illustrates that the “price,” v_{it} , rises faster than it falls. The bid-side spread is larger when prices are rising because dealers can delay recognition of the increase in the prices at which they buy from retail investors. The ask-side spread is larger when prices are falling because the dealers can opportunistically delay recognition of the decrease in the prices at which they sell to some customers, and the measured ask price is an average. For the midpoint to rise faster than it falls, some asymmetry is required between retail customers on the ask side and “wholesale” pricing on the bid side. Our empirical results for both prices and implicit spreads are consistent with these aspects of the setting, and thus provide new evidence that dealers are not behaving in a competitive manner. Retail prices on the bid side fall faster than they rise, and on the ask side they rise faster than they fall. One side of the market, however, is more heavily retail, and so measures of the price (such as the midpoint) rise faster than they fall.

In Part A of the Internet Appendix we evaluate more elaborate and realistic examples and simulations to be sure that our empirical measures of asymmetric movements in prices and dealer spreads capture opportunistic recognition of movements in fundamentals.

III. Price Stickiness and Asymmetric Adjustment

Efficient price discovery in financial markets requires that, when underlying intrinsic values move, prices follow. In this section we ask how prices and volume respond to information, and whether the terms of trade display asymmetric behaviors in rising and falling markets. We show that municipal prices are extremely sticky, and that asymmetries in their sluggish response to news appear to benefit broker-dealers intermediating the trades. In Section IV we further explore whether these asymmetries vary with measures of market power for dealers, and whether they are consistent with formal models of search frictions for customers.

A. News Events, Volume, and Prices

Every month the U.S. government announces macroeconomic indicators such as the Consumer Price Index and Advance Retail Sales. Financial market participants watch these numbers closely to evaluate the state of the economy. The values of tax-exempt bonds, like their taxable counterparts, depend on information about current and future economic outcomes. Where do investors trade across the two linked markets? How quickly are price changes in one market reflected in the other? How much trade is involved in this adjustment process? The tax-exempt market, unlike the Treasury market, involves considerable trade by small investors. How do the terms of trade for retail investors respond to major announcements, and what can this teach us about the role of the broker-dealers intermediating the trades?

Studies on the Treasury market generally find that prices react almost instantaneously to surprises in scheduled macroeconomic announcements, that

the announcements trigger abnormally high volume, and that there is little autocorrelation in returns after the first minute (Ederington and Lee (1993, 1995), Fleming and Remolona (1999), and Balduzzi, Elton, and Green (2001)). Piazzesi (2005) studies the price reaction to the release of FOMC meeting statements and finds that the price response to surprises in these announcements is more sluggish, perhaps because of the qualitative nature of the announcements and their unexpected timing. In all of these studies, however, the time horizon for evaluating price adjustments is measured in minutes. We show that price adjustments in the municipal market are spread out over days.

Given the obvious liquidity advantage of the Treasury market, in which a large volume of trade is concentrated on a small number of bonds, it is not surprising that price discovery takes place there first.¹¹ In fact, we show that virtually all news-related trade takes place in the Treasury market. We find no evidence of significant increases in trading volume in the municipal market; when there is significant abnormal volume, it is generally negative.

To evaluate the impact that news announcements have on volume, we regress total daily volume on indicators for whether the day saw a news announcement of a particular type, interacted with whether the news was good, bad, or neutral. We define news as bad when the surprise in the announcement (the difference between the actual number and the Bloomberg survey median) is unfavorable to the bond market, such as a low unemployment rate or surprisingly high advance retail sales. Across the 1,615 business days in our sample, 459 days have at least one announcement of the indicators listed in Table I, where 233 days have at least one announcement that was good news. Some days with multiple announcements have mixed news. In such cases, the day can simultaneously be a “good news day” for one type of announcement and a “bad news day” for another. There are also cases in which news announcements meet expectations; these are called “neutral news days.” Daily volume is measured as the natural log of par value traded for each of the three trade categories—dealer buys, dealer sales, and interdealer trades.

In Table III we report the results of three such regressions, where the dependent variable is total daily volume of customer buys, customer sales, and interdealer trades, respectively. We aggregate the news events across types of announcements into good news, bad news, and neutral news days based on the average response of Treasuries reported in Table I. Results are very similar when we disaggregate different announcements, as can be seen in Part B of the Internet Appendix. We also include indicators for days following the news events, to allow for a delayed volume response. Because volume in the municipal bond market shows day-of-the-week effects, we add dummy variables to control for this possibility. Finally, we allow the disturbance terms to follow ARMA processes to capture any persistence and we select the lag lengths of the autoregressive and moving-average components according to

¹¹We conduct Granger causality tests to evaluate where price discovery is occurring first. The results are reported in Part B of the Internet Appendix to the paper, and show that the Treasury yield Granger causes the municipal yield, but not the reverse.

Table III
The Volume Response to Macro News

The table reports estimates of the volume response to macroeconomic news. Macroeconomic news events are categorized as *Good News* (at least one good news item on the day), *Bad News* (at least one bad news item), or *Neutral News* (news item released meeting expectations). The dependent variables are the natural logarithms of daily trading volume split by the type of trade (dealer sale, dealer buy, interdealer transaction). We report the coefficients for contemporaneous and lagged news indicators. In addition, we control for day-of-the-week and year fixed effects. We assume that the disturbance terms follow autoregressive processes and select the lag lengths of the autoregressive and moving-average components according to Schwarz's Bayesian information criterion (BIC). The autoregressive specification chosen by the BIC is an ARMA(1,1) for dealer sales and dealer buys and an AR(1) for interdealer transactions. *, **, and *** indicate that the corresponding *p*-values are less than 0.10, 0.05, and 0.01, respectively.

	Dealer Sales	Dealer Buys	Interdealer
Good news _{<i>t</i>}	0.08	0.06	0.09
Good news _{<i>t</i>-1}	-0.07*	-0.11**	-0.15***
Good news _{<i>t</i>-2}	0.01	-0.02	-0.06
Bad news _{<i>t</i>}	0.05	0.05	0.11
Bad news _{<i>t</i>-1}	-0.04	-0.03	-0.06
Bad news _{<i>t</i>-2}	0.01	-0.02	0.02
Neutral news _{<i>t</i>}	0.03	0.03	0.05
Neutral news _{<i>t</i>-1}	-0.05	-0.06	-0.03
Neutral news _{<i>t</i>-2}	0.03	0.05	0.05
BIC	1,753	2,182	2,963
Observations	1,613	1,613	1,613

Schwarz's Bayesian information criterion (BIC). There is no evidence of an increase in volume in response to news releases in the municipal market. The only (marginally) significant coefficients, for good news at a 1-day lag, suggest there is less volume on such days than is usual.

It seems clear, then, that traders wait on news days for price discovery to take place elsewhere, as one might expect when an alternative, highly liquid market is available. Comments from traders are consistent with this view. For example, on October 19, 2006, the *Bond Buyer Online* reported the following:

“The market is sideways for the most part,” a trader in New York said. “However, there still seems to be some interest in bonds. There are some bidders, which is an improvement in this situation from the last week. As far as the data goes, it's funny how we used to really focus and hang on the economic announcements and act accordingly. Now, we don't bother, we just wait to see what the Treasury market does in the wake of the economic data and take our cue from that. Instead of acting, we react.”

Table IV provides descriptive evidence of slow price discovery for individual municipal bonds when there is news that moves interest rates. Most periodic macroeconomic news announcements are released during morning hours, at the times listed in column 3 of Table I. For each announcement day in our

Table IV
The Effects of Macroeconomic News on Yield Spreads

The table documents the effect of macro announcements on municipal bond yields and on spreads between municipal bonds and Treasuries. The yield log-spread is the natural logarithm of the ratio of the maturity-matched Treasury rate and the midpoint yield on the muni bond. Yields are measured in basis points. The explanatory variables capturing the effect of macro news are the standardized surprise component in the macro announcement as described in Table I. Coefficients for additional control variables are tabulated in Part B of the Internet Appendix. These controls include macro announcement dummies that equal one if there is an announcement of the corresponding item on the given day and zero otherwise, the bond and issuer characteristics described in Appendix A, order flow variables (bond-level trading volume and aggregate order imbalances over the past 90 sessions), par sizes, and the daily changes in the size of the trades used to measure muni yields. In addition, we control for calendar year and state of issuance fixed effects. The estimation results are from a cross-sectional regression on the announcement day, and standard errors are adjusted for heteroskedasticity. The sample is restricted to investment-grade bonds. *, **, and *** indicate that the corresponding p -values are less than 0.10, 0.05, and 0.01, respectively.

	Δ Yield Log-Spread	Δ Yield	Δ Treasury Yield
Positive Macro Surprise:			
Advance retail sales	0.41***	0.08	1.90***
Capacity utilization	0.12**	-0.22	0.74***
Nonfarm payrolls	1.20***	1.63***	7.56***
Consumer price index	0.13***	0.40**	0.96***
GDP annualized	0.37***	0.34**	2.26***
Industrial production	-0.02	-0.30	-0.48***
Producer price index	-0.46***	0.08	-2.59***
Consumer confidence	0.12***	0.01	0.82***
Jobless rate	0.17***	-0.14	0.56***
FOMC rate decision	0.03	-0.06	0.04
Negative Macro Surprise:			
Advance retail sales	0.28***	0.31*	1.98***
Capacity utilization	0.46***	-0.01	1.96***
Nonfarm payrolls	0.56***	0.40***	2.88***
Consumer price index	0.41***	0.33	2.27***
GDP annualized	-0.02	0.15	-0.00
Industrial production	-0.36***	0.42**	-1.36***
Producer price index	0.16***	0.14	1.31***
Consumer confidence	-0.03	0.18*	-0.00
Jobless rate	-0.00	-0.18	-0.26***
FOMC rate decision	-0.05	0.19	0.04*
Observations	205,428	207,210	207,210
R^2	0.11	0.10	0.15

sample, we consider all bonds with sufficient trades after the first announcement on that day and the day before to allow us to compute midpoint yields as described in Section I. Using these midpoint yields for each day, we compute the announcement effect as the change in yield. Positive macro surprise variables are, again, the standardized announcement for a particular news item on days that the announcement exceeds the median forecast, and zero on all other days.

The median forecast is from survey data compiled by Bloomberg. The negative surprise variables are constructed analogously.¹² We conduct a cross-sectional regression of the daily change in the yield against the surprise variables and the set of controls described in Appendix A. We have 775 announcement events on 459 days, and an average of 452 bonds with sufficient trades on each of these days, for a total of 207 thousand observations. The table reports the coefficient estimates for the various news surprise variables. The table shows that municipal bonds underreact to macroeconomic news. The economic magnitudes of the coefficients are small compared to those for Treasuries, and the coefficients are often statistically insignificant, again in contrast to Treasuries. Yield spreads react dramatically to macroeconomic news, as the first column in the table shows because Treasury rates respond quickly (third column) while municipal rates are extremely sluggish (second column).

Note that our data are transaction prices, not posted quotes. Thus, we are observing lagged adjustment in the prices at which bonds are actually being exchanged. Some of this slow response appears to be a failure of prices to move at all. In Figure 1 we plot the distribution of percentage changes in the yields of municipal bonds and in the yields of maturity-matched Treasuries. For every bond and every day in the sample, we check to see if we can construct a midpoint for the inside effective spread as described in Section I. We then attempt to match the resulting bond yield with a yield on the same bond 1 day later, 5 days later, and 10 days later. We pool all the bond-days for which these pairs can be constructed, and plot kernel estimators for the frequency distribution of the percentage changes in yield. The distribution for the individual municipals is initially more peaked, and gradually assumes a shape similar to that evident for the Treasuries. Much of the more concentrated mass for the municipals, however, is attributable to a spike at exactly zero. That is, transaction prices are simply not moving over time, even as the maturity-matched Treasury moves. In many cases, it appears that price adjustment is not only sluggish or slow—it is not occurring at all.

B. Dealer Markups on Round-Trip Trades

We have seen that price adjustment is delayed in the municipal market. We next ask whether the terms of trade display asymmetric behaviors in rising and falling markets, and across retail versus institutional investors. Do prices rise faster than they fall?

The municipal market involves trade in a large number of bonds, but trade is sparse on the time-series dimension. Since price discovery takes place over time, this presents challenges that we address through several different strategies. As a first step, we follow Green, Hollifield, and Schürhoff (2007a) and measure the profits that dealers earn on round-trip transactions, initiated by

¹²Note that “positive” and “negative” surprises do not correspond to “good news” and “bad news” surprises, but simply to the sign of the deviation from the median forecast. A positive surprise for advanced retail sales, for example, is “bad news” for bond prices, as shown in Table I.

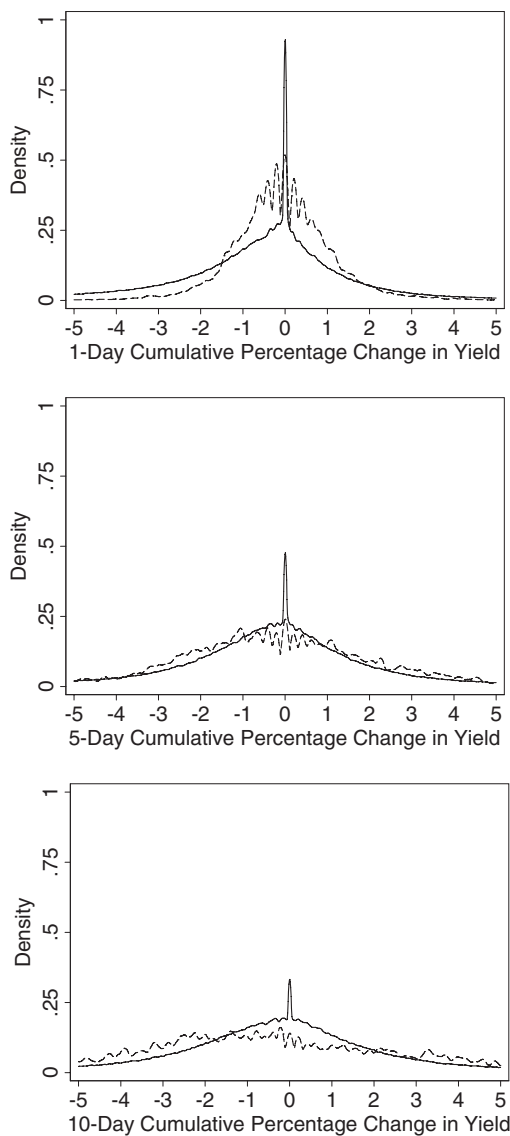


Figure 1. Distribution of changes in yields on munis and Treasuries. The different panels plot the distribution of daily percentage changes in yields on municipal bond transactions (solid line) and the corresponding distribution of percentage changes in Treasury rates (dashed line) at increasingly long horizons. The muni yields are daily midpoints for each bond issue traded. The Treasury rates are maturity-matched.

a purchase of bonds from a customer. The MSRB data identify trades as dealer purchases from customers, dealer sales to customers, or interdealer trades. The data do not reveal the identities of specific dealers or customers. Since trade is relatively infrequent, however, it is often evident that specific transactions are the legs of transactions intermediated by the same dealer, for example, when

a purchase from a customer is followed by a sale of the same bond in the same par amount with no intervening transactions, or when a purchase is followed by two sales of the same bond that sum to the initial par value.

We match purchases from customers with subsequent sales as follows. Trades that are customer sales to a dealer immediately followed by a customer buy from a dealer at the same par amount are identified as a round-trip pair. The matched trades are then filtered according to the following criteria. If the yield on either leg is missing, the yield is greater than 20%, the price on either side is smaller than \$20 or greater than \$150, the par value is less than \$5,000, or the percentage markup is greater than 10% or less than -10%, then the record is deleted. Since we are interested in price responses to market movements, and market prices are available only at a daily frequency, we further eliminate round-trip pairs that happened on the same day and pairs for which the legs of an apparent trade are more than 10 calendar days apart. There are 2,156,188 pairs left after filtering. Among them, 69,528 (3%) pairs are large institutional size (Par \geq \$250 thousand), 446,678 (21%) pairs are medium size (Par \in [\$50, \$250) thousand), and 1,639,982 (76%) pairs are small retail size.

Dealers in the seasoned municipal market do not routinely maintain inventories in the large numbers of bonds they trade. They also have no mechanism for shorting individual bonds. Thus, their role as liquidity provider is typically triggered by a purchase from a customer, followed by a search for buyers. Customers needing to buy are more likely to be serviced by the dealer searching for and identifying a potential seller, negotiating the terms, and executing the two sides of the trade simultaneously. This is evident in our data set, which contains over 3.5 million round-trip purchases from customers followed by a sale of the same bond in the same par amount, out of which 1.3 million are completed on the same day. Over 70% of the two legs are separated in time by more than 2 minutes and less than 5 days. In contrast, if we try to match transactions in the reverse order, sales to customers followed by a purchase of the same bond in the same amount, we can match only one million pairs (many of which involve trades that can also be matched as round trips). Of these, 19.4% involve trades where the sale and purchase are within 2 minutes of each other, suggesting the trade was prearranged, while 48.1% have two legs more than 5 days apart, suggesting it is unlikely the two sides are moving through the same dealer. Part A of the Internet Appendix provides details on the distribution of time between trades in different directions.

Once transactions are matched, a simple measure of the profits a dealer (or dealers, if some trades are incorrectly “matched”) makes on a given sequence of trades is the percentage markup over the purchase price—the difference between the proceeds from the dealer’s sales to customers and the cost of buying the bonds, divided by that cost.

Other authors, such as Harris and Piwowar (2006) and Green, Hollifield, and Schürhoff (2007a), investigate the cross-sectional determinants of dealer trading profits, but not whether markups differ when prices are rising versus falling. If prices rise faster than they fall, as in markets for retail goods, then the

markup should increase during market rallies by more than it falls when prices are decreasing. Dealers can opportunistically delay responding to price drops, thus preserving their markup, while immediately recognizing price increases, increasing their profits on a trade. This is exactly what we find at the retail level. In fact, markups increase when prices rise but are insensitive to price movements when prices fall. Institutional investors are more sophisticated shoppers for the services of intermediaries. They are informed about market conditions, and have repeated interactions with dealers. This puts them in a superior bargaining position. Institutionally sized trades do not show this asymmetry.

To measure the underlying price movements, we use the Lehman Brothers Long Term Municipal Price Index to proxy for the market price level. The Lehman Brothers Long Term Municipal Index is a benchmark index that includes investment-grade, tax-exempt, and fixed-rate bonds with long-term maturities (greater than 22 years) selected from issues larger than \$50 million.

We regress the markup on round-trip transaction i against the change in the index over the period between the initial purchase and the final sale. The model in equation (1) separates rising and falling markets, and allows the coefficients to differ:

$$\text{Markup}_i = \beta^+(r_t^{\text{idx}})^+ + \beta^-(r_t^{\text{idx}})^- + \gamma'x_{it} + \varepsilon_i, \quad (1)$$

where Markup_i is expressed as a percent and x_{it} is a set of explanatory variables. The variables $(r_t^{\text{idx}})^+$ and $(r_t^{\text{idx}})^-$ are the returns in the underlying index (computed as log changes) when that return is positive and negative, respectively. Control variables are the number of days between the legs of the transaction, the log par value of the trade, an indicator for whether the markup is rounded (to at least 1/8th of a dollar), and an intercept term. All of these variables are known to be related to the size of dealer markups (see, e.g., Green, Hollifield, and Schürhoff (2007a) and Li (2007)).

In Table V we report the results of estimating this model separately for small-medium- and large-sized trades using a panel regression with bond-specific fixed effects. Standard errors in this and subsequent tables are adjusted to account for two-way clustering, by issuer and by day (the details are presented in Part C of the Internet Appendix). We expect institutions, which trade in larger sizes, to be more aware of market conditions, and hence to be more effective bargainers. We also separate trades where the dealer holds the bonds in inventory for more than five days from those of more typical duration. As dealers begin to get desperate to unload their trades in falling markets, the terms of trade may become more responsive to price movements. The p -values in the table test the restriction that $\beta^- \geq \beta^+$.

As the table shows, when prices rise, markups rise, regardless of trade size or the length of time the dealer holds the position. When prices fall, markups also fall (there is a positive coefficient on $(r_t^{\text{idx}})^-$). For small- and medium-sized trades of routine duration, however, this coefficient is insignificant or only marginally significant (keep in mind the very large sample), and it is much

Table V
Markups in Rising and Falling Markets

The table documents percentage markups on round-trip transactions in rising and falling markets. The sample is split into six categories depending on the amount of time between two legs of the round-trip (ΔT) and the par size of the order. Small, Medium, and Large correspond to par sizes smaller than 50K, between 50K and 250K, and greater than or equal to 250K. The variables $(r_t^{idx})^-$ and $(r_t^{idx})^+$ are the returns in the underlying index (computed as log changes) when they are negative and positive, respectively. Control variables are the number of days between the legs of the transaction, the log par value of the trade, and an indicator for whether the markup is rounded (to at least 1/8th of a dollar). The estimation results are from a panel regression with bond-specific fixed effects. In parentheses we report cluster-robust t -statistics for $H_0 : \beta = 0$. Standard errors are adjusted to account for two-way clustering on bond issuer and calendar time. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006. Outlier observations for which the absolute value of round-trip markups is greater than \$10 per \$100 par value, and those for which the absolute value of the markups is greater than 10%, are dropped.

Markup (%)	$\Delta T \in [1, 5]$			$\Delta T \in [6, 10]$		
	Small	Medium	Large	Small	Medium	Large
	Coef. (t)	Coef. (t)	Coef. (t)	Coef. (t)	Coef. (t)	Coef. (t)
Rising market: $(r_t^{idx})^+$	0.19 (6.05)	0.18 (4.81)	0.16 (1.82)	0.26 (10.19)	0.26 (4.26)	0.42 (1.79)
Falling market: $(r_t^{idx})^-$	0.03 (1.32)	0.06 (1.82)	0.20 (2.24)	0.26 (12.81)	0.30 (5.33)	0.36 (1.62)
Control variables:						
Time between legs	0.01 (2.51)	0.01 (2.41)	-0.00 (-0.06)	-0.01 (-6.72)	-0.01 (-2.27)	-0.01 (-0.46)
Par size (log)	-0.44 (-82.60)	-0.33 (-32.35)	-0.22 (-7.94)	-0.36 (-29.58)	-0.34 (-5.99)	-0.16 (-1.48)
Markup rounded	-0.28 (-38.25)	-0.05 (-2.38)	-0.03 (-0.20)	-0.35 (-15.31)	-0.08 (-0.51)	0.31 (0.36)
p -value for $H_0 : \beta^- \geq \beta^+$	<.001	0.02	0.60	0.43	0.66	0.43
Observations	1,291,624	351,897	52,377	348,358	94,781	17,151
Adj. R^2	0.30	0.43	0.52	0.24	0.43	0.51

smaller in magnitude than the corresponding coefficient for rising markets. Dealers selling to retail customers raise their prices sharply when the index rises but do not drop their prices comparably when the index falls. This stands in contrast to how the intermediaries deal with institutional investors, where the asymmetry is both statistically and economically insignificant, and if anything goes in the opposite direction. Institutional investors may, therefore, have a small bargaining advantage when prices are falling and dealers hold bonds in their inventories.

C. Implicit Half-Spreads and Asymmetric Price Adjustments

The analysis of round-trip profits for dealers imposes an implicit ordering on transactions and excludes a great deal of data. It examines only trades that

can be matched as likely round-trips. Because we need index data to measure movements in fundamentals, we can use only round-trip trades that carry over across days.

In the tests that follow we aggregate data for specific bonds at the daily level. We compute daily bid, ask, and midpoint prices for each bond and for different transaction size categories. That is, we construct a daily proxy for the implicit spread conditional on size, rather than estimate the bid-ask spread through a time-series regression on a trade-direction indicator, as in Schultz (2001) and others.

Our proxy for the ask (bid) price is the average transaction price associated with all customer purchases (sales) within a given size category on a given day. The common value of the bond is captured by the midpoint of the highest bid and the lowest ask transaction on a given day if both buys and sales occur. If not, we use the average price on all interdealer transactions.

First, we analyze how bid-ask spreads and half-spreads respond to changes in the midpoint. Let the average price at which dealers sell bond i to customers on day t be denoted p_{it}^{ask} , and let p_{it}^{bid} denote the average price at which the bond is bought from customers. Assume bond i has common value v_{it} at time t . We measure the ask-side and bid-side half-spreads by $sp_{it} = (p_{it}^{ask} - v_{it})/v_{it}$ and $sp_{it} = (v_{it} - p_{it}^{bid})/v_{it}$, respectively. The following model allows for the possibility that transaction prices react asymmetrically to gains and drops in the bond's fundamental value and is estimated using panel data methods:

$$sp_{it} = \beta^+(\Delta v_{it})^+ + \beta^-(\Delta v_{it})^- + \gamma'x_{it} + \alpha_i + \varepsilon_{it}, \quad (2)$$

where Δv_{it} is expressed as a percent, x_{it} is a set of explanatory variables, and $\varepsilon_{it} = \rho\varepsilon_{it-1} + \xi_{it}$ is an error term, which is allowed to be autocorrelated. The bond-specific random effect is α_i . For a bond-day to be included in the sample for these tests we must have at least one trade on each side of the market in that bond to construct the midpoint, or we must have an interdealer trade to proxy for the midpoint. In the latter case, we might have one half-spread but not the other, with resulting differences in the sample size.

Suppose dealers can opportunistically delay their response to changes in fundamentals due to search costs for investors. Then we would expect the bid-side spreads to increase when the midpoint rises. In the extreme, for example, if dealers could hold the price at which they buy fixed, the bid-side spread would rise one-to-one with the midpoint. The reverse holds for the ask-side spread when prices fall. Dealers would wish to delay dropping the price at which they could sell to investors. If competition were less than perfect, they might succeed in this, particularly if the source of imperfect competition lies with costly search for the best price. In contrast, dealers would be quick to alter the price at which they buy (sell) when prices fall (rise). We would then expect the bid-side (ask-side) spread to remain stable when the midpoint falls (rises).

Table VI reports results on the effect of movements in the bond's fundamental value, estimated separately on sell-side and buy-side half-spreads. We

Table VI
The Effect of Movements in Fundamentals on Bid-Ask Spreads and Half-Spreads

The table documents the asymmetric effect of movements in a municipal bond's common value on bid-ask spreads and half-spreads. The explanatory variables $(\Delta v_{it})^+$ and $(\Delta v_{it})^-$ are the positive and negative components, respectively, of the change in the midpoint bond price, expressed in percent. Coefficients for additional control variables are tabulated in Part B of the Internet Appendix, including the bond and issuer characteristics and order flow variables described in Appendix A, the maturity-matched interest rate at the time of trade, and measures capturing trade size and the daily changes in the size of the trades used to measure bids, asks, and midpoints. The estimation results are from a panel regression with bond-specific random effects. Small trades are less than \$50K in par value, medium are between \$50K and \$250K, and large trades exceed \$250K. In parentheses we report cluster-robust t -statistics for $H_0 : \beta = 0$. Standard errors are adjusted to account for two-way clustering on bond issuer and calendar time. P -values for the hypothesis $H_0 : \beta^+ \leq -\beta^-$ are one minus the p -values reported for each column. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006, aggregated at the daily frequency.

	Bid-Ask Spread (Ask-Bid)/Mid			Ask Half-Spread (Ask-Mid)/Mid			Bid Half-Spread (Mid-Bid)/Mid		
	Small	Medium	Large	Small	Medium	Large	Small	Medium	Large
Rising value: $(\Delta v_{it})^+$	Coef. (t) 2.62 (4.20)	Coef. (t) 10.84 (8.59)	Coef. (t) 15.41 (11.74)	Coef. (t) 0.61 (1.84)	Coef. (t) 2.90 (5.09)	Coef. (t) 5.58 (8.58)	Coef. (t) 8.55 (20.64)	Coef. (t) 8.30 (17.41)	Coef. (t) 12.18 (24.44)
Falling value: $(\Delta v_{it})^-$	Coef. (t) -29.69 (-36.11)	Coef. (t) -30.35 (-12.43)	Coef. (t) -24.63 (-13.18)	Coef. (t) -30.30 (-62.03)	Coef. (t) -46.84 (-39.14)	Coef. (t) -41.91 (-35.10)	Coef. (t) -7.23 (-15.84)	Coef. (t) -2.78 (-3.25)	Coef. (t) -5.67 (-7.25)
p -value for $H_0 : \beta^+ \geq -\beta^-$	<.001	<.001	<.001	<.001	<.001	<.001	.999	.999	.999
Observations	193,515	36,620	44,727	520,813	203,125	190,315	251,837	83,990	85,118
R^2	0.25	0.37	0.33	0.34	0.39	0.36	0.17	0.24	0.24

also estimate the model for the bid-ask spread (average customer buying price minus average customer selling price on a given day, normalized by the midpoint). Results in the table are stratified by transaction size. Large trades involve par values in excess of \$250 thousand, medium sizes are between \$50 thousand and \$250 thousand, and small trades are less than \$50 thousand.

The coefficient estimates suggest, first, that larger price movements increase the effective bid-ask spread. In the first three columns, the reported coefficients are positive when the change in the midpoint is positive ($\beta^+ > 0$), and negative when the change in the midpoint is negative ($\beta^- < 0$). Together, these imply an expected increase in the bid-ask spread in response to changes in the midpoint in either direction. The associated coefficients for the half-spreads display asymmetric magnitudes across price drops and increases. P -values for tests of equality for β^+ and $-\beta^-$ are virtually zero for every subsample in the table. The largest coefficients in absolute value are associated with the mid-price less the bid when the midpoint is rising, and the ask less the mid-price when the midpoint is falling. The ask-side effective spread increases dramatically when prices fall, and the bid-side effective spread rises the most when prices rise. This pattern occurs consistently across transaction size categories. This behavior is consistent with dealers responding relatively quickly to movements in fundamentals when it benefits them to do so, while resisting the competitive pressure to change prices quickly when doing so erodes their markups. The largest coefficients are for the bid-ask spread and the ask-side half-spread when prices fall. This suggests that dealers are not revising the prices at which they sell to customers when prices are dropping. Prices rise faster than they fall on the ask side of the market, and they fall faster than they rise on the bid side. The asymmetric behavior is most pronounced on the ask side, where dealers are generally trading in smaller quantities.

We also estimate a partial adjustment model for the effective bid-ask spreads, for ask-side half-spreads (i.e., ask minus midpoint, normalized by midpoint), and for bid-side half-spreads (i.e., midpoint minus bid, normalized by midpoint) using panel data methods:

$$\Delta sp_{it} = \begin{cases} \delta^+ (sp_{it}^* - sp_{it-1}), & \text{if } sp_{it}^* \geq sp_{it-1}, \\ \delta^- (sp_{it}^* - sp_{it-1}), & \text{if } sp_{it}^* < sp_{it-1}. \end{cases} \quad (3)$$

The term in parentheses in (3) captures the deviation from equilibrium, ϵ . The parameters δ^- and δ^+ measure the adjustment speed since $\delta^- = \lim_{\epsilon \uparrow 0} \frac{\partial}{\partial \epsilon} \Delta sp$ and $\delta^+ = \lim_{\epsilon \downarrow 0} \frac{\partial}{\partial \epsilon} \Delta sp$. The model postulates that the change in the dealer spread depends on the distance between its previous value and a proxy for the “equilibrium” value, sp_{it}^* . We model the latent equilibrium value as $sp_{it}^* = \beta' x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable characteristics of bond i and the market environment on date t , α_i is a bond-specific random effect, and the residual ε_{it} is normally distributed. Bond characteristics include maturity, coupon, and various credit quality measures. Market conditions include lagged volume and measures of average trade size. Past research shows these variables influence

trading costs. The speed of adjustment parameters, δ^- and δ^+ , can differ based on whether the dealer spread is above or below its latent value. Part C of the Internet Appendix provides details on how we construct the maximum likelihood estimator of (3).

Table VII reports maximum likelihood estimates for the parameters in (3). The results suggest that bid-ask spreads and half-spreads adjust more quickly when they are too narrow than when they are too wide. The numbers in parentheses below the coefficient estimates report t -statistics for tests of whether the adjustment coefficients are equal to one. The adjustment coefficients δ^+ , which measure speed of adjustment when bid-ask spreads and half-spreads are below the latent “equilibrium” values, are economically close to unity, and have relatively small t -statistics associated with the deviation from unity for the bid-ask spread and the bid-side half-spread. The second row from the bottom of the table reports tests for asymmetric adjustment. The null hypothesis is that $\delta^- = \delta^+$. The adjustment coefficients when the bid-ask spreads and half-spreads are larger than the latent values are in every case smaller, significantly so in both economic and statistical terms.

D. Asymmetric Yield Spread Dynamics

Next, we provide evidence on how quickly spreads between municipal and Treasury yields adjust to changes in fundamentals. Let $s_{it} = \ln(y_{it}^T) - \ln(y_{it}^M)$ denote the log-yield spread between municipal bond i , evaluated at the midpoint for day t , and a maturity-matched Treasury. The model we estimate is

$$\Delta s_{it} = \begin{cases} \delta^+(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* \geq s_{it-1}, \\ \delta^-(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* < s_{it-1}. \end{cases} \quad (4)$$

Again, the latent equilibrium value is modeled as $s_{it}^* = \beta'x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable controls, α_i is a bond-specific random effect, and ε_{it} is normally distributed.

Note that when $s_{it}^* > s_{it-1}$, the previous day’s yield spread is “too low” relative to the current predicted value. Alternatively, the Treasury yield is “too low” relative to the municipal yield. Municipal yields should fall relative to Treasury yields, or municipal prices should rise. Therefore, we interpret δ^+ as measuring the speed of adjustment when prices are rising relative to Treasuries, and δ^- as the speed of adjustment when prices are falling.

Table VIII reports the results, stratified by bond rating and by year. Assuming Treasury yields adjust quickly and appropriately to new information, it is evident that municipal yields adjust sluggishly. The partial adjustment coefficients are generally significantly different from one. In every case, the partial adjustment coefficient for rising markets, δ^+ , is larger than that for falling markets, δ^- . Prices rise faster than they fall in the municipal bond market. The asymmetry does not tend to decrease in any obvious way over time, and is evident for all ratings categories. Thus, while dealers appear to be able to opportunistically time the recognition of price movements on both

Table VII
Bid-Ask Spread Dynamics

The table reports parameter estimates for an asymmetric- and partial-adjustment model of effective bid-ask spreads. Let sp_{it} be the effective spread in bond i at time t , and denote by sp_{it}^* the unobserved equilibrium value. The specification we estimate is the following:

$$\Delta sp_{it} = \begin{cases} \delta^+ (sp_{it}^* - sp_{it-1}), & \text{if } sp_{it}^* \geq sp_{it-1}, \\ \delta^- (sp_{it}^* - sp_{it-1}), & \text{if } sp_{it}^* < sp_{it-1}. \end{cases}$$

This specification allows for asymmetric partial adjustment. The parameters δ^- and δ^+ measure the adjustment speed since $\delta^- = \lim_{\epsilon \uparrow 0} f'(\epsilon)$ and $\delta^+ = \lim_{\epsilon \downarrow 0} f'(\epsilon)$. Further, we assume $sp_{it}^* = \beta x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable characteristics of bond i at time t , α_i is a bond-specific random effect, and the residual ε_{it} is normally distributed. The specification for sp_{it}^* includes additional explanatory variables (dummy variables for each rating category and issuer type variables as in Table VI) and control variables for the par sizes and the daily changes in the size of the trades used to measure bids, asks, and midpoints. Coefficients for these controls are tabulated in Part B of the Internet Appendix. In parentheses we report cluster-robust t -statistics corresponding to $H_0 : \delta = 1$ and $H_0 : \beta = 0$, respectively. Standard errors are adjusted to account for two-way clustering on bond issuer and calendar time. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006, aggregated at the daily frequency.

	Bid-Ask Spread (Ask-Bid)/Mid	Ask Half-Spread (Ask-Mid)/Mid	Bid Half-Spread (Mid-Bid)/Mid
Speed of Adjustment:	Coef. (t vs. 1)	Coef. (t vs. 1)	Coef. (t vs. 1)
δ^+	0.98 (1.22)	0.92 (10.15)	0.98 (2.10)
δ^-	0.91 (7.10)	0.86 (19.61)	0.89 (11.28)
Explanatory variables β :	Coef. (t vs. 0)	Coef. (t vs. 0)	Coef. (t vs. 0)}
Maturity	0.39 (7.86)	0.16 (8.79)	0.18 (8.39)
Coupon	-0.40 (-1.07)	-0.39 (-2.94)	0.24 (1.69)
Zero coupon bond	4.97 (2.52)	1.39 (2.07)	5.18 (6.88)
Callable bond	3.84 (15.14)	1.97 (18.82)	1.92 (17.94)
Insured bond	-0.30 (-0.88)	-0.27 (-2.71)	-0.59 (-4.15)
Rating BB-B	4.99 (7.65)	3.53 (11.65)	2.49 (8.39)
Revenue bond	0.65 (3.46)	0.22 (3.16)	0.46 (5.20)
Issue size	0.14 (0.30)	-0.34 (-4.26)	0.14 (0.72)
Volume [$t-90, t-1$]	-0.02 (-2.64)	0.01 (2.82)	-0.02 (-7.42)
p -value for $H_0 : \delta^- \geq \delta^+$	<.001	<.001	<.001
Observations	187,928	449,742	264,855

Table VIII
Yield Spread Dynamics

The table reports parameter estimates for the asymmetric and partial adjustment model. Let s_{it} be the yield log-spread of the maturity-matched Treasury rate s_{it}^T over the yield s_{it}^M on muni bond i at time t , and denote by s_{it}^* the unobserved equilibrium value. The specification we estimate is the following:

$$\Delta s_{it} = \begin{cases} \delta^+ (s_{it}^* - s_{it-1}), & \text{if } s_{it}^* \geq s_{it-1}, \\ \delta^- (s_{it}^* - s_{it-1}), & \text{if } s_{it}^* < s_{it-1}. \end{cases}$$

This specification allows for asymmetric partial adjustment. The parameters δ^- and δ^+ measure the adjustment speed since $\delta^- = \lim_{\epsilon \uparrow 0} f'(\epsilon)$ and $\delta^+ = \lim_{\epsilon \downarrow 0} f'(\epsilon)$. Further, we assume $s_{it}^* = \beta x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable characteristics of bond i at time t , α_i is a bond-specific random effect, and the residual ε_{it} is normally distributed. The specification for s_{it}^* includes the following control variables, for which results are tabulated in Part B of the Internet Appendix: bond issue and issuer characteristics, order flow variables, macro variables as described in Appendix A, dummies for the calendar year, and control variables for the par sizes and the daily changes in the size of the trades used to measure municipal yields. In parentheses we report cluster-robust t -statistics corresponding to $H_0 : \delta = 1$. Standard errors are adjusted to account for two-way clustering on bond issuer and calendar time. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006, aggregated at the daily frequency.

Panel A: Split by Rating							
	All	Unrated	AAA	AA	A-BBB	BB-B	CCC-D
Speed of Adjustment:	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.
	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)
δ^+	0.92 (3.72)	0.98 (1.34)	0.92 (3.04)	0.95 (1.79)	0.72 (6.31)	0.80 (2.78)	0.62 (4.66)
δ^-	0.74 (17.35)	0.70 (21.95)	0.80 (9.87)	0.88 (5.70)	0.63 (9.63)	0.73 (3.28)	0.59 (4.53)
p -value for $H_0 : \delta^- \geq \delta^+$	<.001	<.001	<.001	<.001	<.001	<.001	0.001
Observations	801,564	252,920	393,735	64,105	73,473	10,704	6,570
Panel B: Split by Year							
	2000	2001	2002	2003	2004	2005	2006
Speed of Adjustment:	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.	Coef.
	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)	(<i>t</i> vs. 1)
δ^+	0.97 (0.54)	0.90 (1.47)	0.98 (0.73)	0.88 (2.22)	0.97 (0.32)	0.93 (1.62)	0.91 (1.90)
δ^-	0.82 (4.26)	0.75 (3.86)	0.75 (8.38)	0.67 (8.44)	0.77 (3.34)	0.76 (6.08)	0.80 (4.00)
p -value for $H_0 : \delta^- \geq \delta^+$	<.001	<.001	<.001	<.001	<.001	<.001	<.001
Observations	83,662	105,271	117,901	125,819	128,765	105,456	134,543

sides of the market, the asymmetry is stronger on the ask side, where more of their customers are retail. Thus, the midpoint inherits the asymmetry associated with the ask side. On average, dealers are “buying wholesale and selling retail.”

IV. Competitive Conditions, Search Costs, and Asymmetric Adjustment

It is difficult to reconcile prices rising faster than they fall with perfect competition. Since there are many broker-dealers in the municipal market, frictions of some sort must be present that keep customers from quickly locating the bonds at the best price. Search costs due to a lack of market transparency may be a source of local market power for broker-dealers. Models of the rockets-and-feathers phenomenon in consumer markets rely on search costs and price dispersion to generate these dynamics. In this section we examine the effects of variation in market conditions, and the implications of models of asymmetric price adjustment based on search.

A. Variation in Competitive Conditions across States

We can interpret asymmetric price adjustments as evidence of market frictions and imperfect competition with more confidence if we can show that the asymmetries vary with competitive conditions. State-level tax treatment of municipal bond interest is a natural place to look for such variation. Most states exempt interest from state income tax only for municipal bonds issued within that state. Investors in such states therefore have a preference for in-state bonds. Past research such as Kidwell, Koch, and Stock (1984) shows that municipal markets are segmented across states, and that the preferential treatment of in-state bonds results in lower yields for bonds issued in those states. There are 37 such states, which we refer to as “segmented states.”¹³ Investors in states with no state income tax or in states taxing both in-state and out-of-state bonds participate in a national market, as do investors in bonds issued in U.S. territories, which are exempt from state taxes due to federal law. Two states, Utah and Indiana, exempt bonds from all states where their own bonds are not taxed. Other factors equal, we would expect less competition among dealers in markets for bonds issued in segmented states.

To test this hypothesis, we allow the coefficients δ^+ and δ^- in the partial adjustment model (4) to depend on whether the state is segmented and on other variables that we expect to be associated with local market power for broker-dealers. The model we estimate is

$$\Delta s_{it} = \begin{cases} \delta^+(z_{it})(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* \geq s_{it-1}, \\ \delta^-(z_{it})(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* < s_{it-1}, \end{cases} \quad (5)$$

¹³Segmented states are AL, AR, AZ, CA, CO, CT, DE, GA, HI, ID, KS, KY, LA, MA, MD, ME, MI, MN, MO, MS, MT, NC, ND, NE, NH, NJ, NM, NY, OH, OR, PA, RI, SC, TN, VA, VT, and WV.

where the speed of adjustment parameters (δ^+ , δ^-) are log-linear functions of the explanatory variables z_{it}^j , $j = 1, \dots, J$:

$$\begin{aligned}\delta^+(z_{it}) &= \delta_0^+ \prod_{j=1}^J e^{\delta_j^+ z_{it}^j}, \\ \delta^-(z_{it}) &= \delta_0^- \prod_{j=1}^J e^{\delta_j^- z_{it}^j}.\end{aligned}\tag{6}$$

As before, we model the latent value for the yield spread as $s_{it}^* = \beta' x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable characteristics of bond i at time t , α_i is a bond-specific random effect, and residual ε_{it} is normally distributed.

The explanatory variables z_{it} for bond i at time t include a dummy for issuance in a segmented state, the state's income tax rate interacted with the segmentation dummy, a measure of underwriter concentration, the size of the state's municipal market, and the amount of retail volume. The measure of underwriter concentration is the Gini coefficient for the state of issuance, where we use SDC Platinum's data on the lead underwriter for each issue to calculate the underwriter's market share of new issues in the state since January of 2000. The local market size is measured by the state's new issues over the sample period as a fraction of the total volume of new issues. The extent of retail participation in secondary market trading is measured by the fraction of trades less than \$50 thousand in par value in the corresponding state. The explanatory variables x_{it} include the variables z_{it} and all order flow variables and bond characteristics included in the basic specification.¹⁴

Recall that δ^+ measures the speed of adjustment when municipal yields must fall relative to Treasuries, so that municipal prices must rise, while δ^- measures the adjustment speed when municipal prices fall relative to Treasuries. The first row of Table IX shows that tax segmentation has no effect on the adjustment speed when prices are rising (column 1), but reduces the adjustment speed when prices are falling (column 2). Tax segmentation thus increases the difference between δ^+ and δ^- (column 3). The adjustment speed is more asymmetric in tax-segmented states, consistent with the view that tax segmentation offers dealers greater opportunity to exercise local market power. Prices rise faster than they fall *more* in states where investors cannot shop in a national market.

While state of issuance is exogenous, states' decisions to tax income of state residents or to exempt in-state bonds are not. The segmented states include the most populous states with the highest tax rates, with the largest numbers of high-net-worth residents, and with the largest bond issues that attract high levels of retail participation. The remaining rows of Table IX attempt to control for these factors. The asymmetry between the speed with which prices rise and

¹⁴Part B of the Internet Appendix reports estimates of the partial adjustment model in (4) on a state-by-state basis, and we report t -tests of the hypothesis that the difference between δ^+ and δ^- is larger for the segmented states as a group. The findings are similar to those reported in the first row of Table IX.

Table IX

Longitudinal Differences in Speeds of Adjustment and Asymmetry

The table reports parameter estimates for the asymmetric and partial adjustment model. Let s_{it} be the yield log-spread of the maturity-matched Treasury rate s_{it}^T over the yield s_{it}^M on municipal bond i at time t , and denote by s_{it}^* the unobserved equilibrium value. The specification we estimate is the following:

$$\Delta s_{it} = \begin{cases} \delta^+(z_{it})(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* \geq s_{it-1}, \\ \delta^-(z_{it})(s_{it}^* - s_{it-1}), & \text{if } s_{it}^* < s_{it-1}, \end{cases}$$

where the speed of adjustment parameters (δ^+, δ^-) are log-linear functions of the explanatory variables $z_{it}^j, j = 1, \dots, J$:

$$\delta^+(z_{it}) = \delta_0^+ \prod_{j=1}^J e^{\delta_j^+ z_{it}^j},$$

$$\delta^-(z_{it}) = \delta_0^- \prod_{j=1}^J e^{\delta_j^- z_{it}^j}.$$

Further, we assume $s_{it}^* = \beta x_{it} + \alpha_i + \varepsilon_{it}$, where x_{it} are observable characteristics of bond i at time t , α_i is a bond-specific random effect, and the residual ε_{it} is normally distributed. The explanatory variables z_{it} of bond i at time t include a dummy for issuance in a segmented state, the state's income tax rate interacted with the segmentation dummy, the Gini coefficient of underwriter concentration in the state's primary market, the local market size proxied by the state's new issues over the sample period as a fraction of total new issuance volume, and the extent of retail participation in secondary market trading proxied by the fraction of trades less than \$50K in par value in the corresponding state. States classified as segmented are all states in which state tax is levied only on bonds issued out of state. Segmented states are AL, AR, AZ, CA, CO, CT, DE, GA, HI, ID, KS, KY, LA, MA, MD, ME, MI, MN, MO, MS, MT, NC, ND, NE, NH, NJ, NM, NY, OH, OR, PA, RI, SC, TN, VA, VT, and WV. Bonds of segmented states are traded mostly in the state of issuance due to local tax advantages. The bonds of the remaining 16 states and territories trade in nonsegmented markets. The explanatory variables x_{it} include the variables z_{it} and all order flow variables and bond characteristics included in the basic specification in Table VIII. In parentheses we report χ^2 -statistics for Wald tests of $H_0 : \delta_j^+ = 0, \delta_j^- = 0$, and $\delta_j^+ = \delta_j^-$, respectively. Critical values for significance at 90%, 95%, 99%, and 99.9% levels are 2.71, 3.84, 6.63, and 10.83. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006, aggregated at the daily frequency.

	δ_j^+	δ_j^-	$\delta_j^+ - \delta_j^-$
	Coef. (χ^2)	Coef. (χ^2)	Coef. (χ^2)
Segmented state (0/1)	0.01 (0.24)	-0.05 (2.10)	0.07 (8.61)
State tax rate (%) \times Segmented state (0/1)	0.43 (1.11)	1.55 (7.79)	-1.11 (9.49)
Primary market concentration (Gini)	0.03 (0.07)	-0.16 (1.28)	0.19 (4.15)
Local market size (% of new issues)	-0.20 (0.99)	0.29 (1.04)	-0.49 (6.79)
Retail participation (1/avg ask size)	-0.28 (7.59)	-0.28 (4.75)	-0.01 (0.01)

fall decreases with the state income tax in states that do not tax municipal bond interest. The asymmetry is also lower in states with a larger municipal market. By these measures, big, high-tax states with more sophisticated high-net-worth investors, such as California, New York, and Pennsylvania, appear to have more competitive and efficient municipal bond markets. Further, the speed of adjustment in rising markets increases, while that for falling markets decreases, in states where underwriting business is more concentrated across dealers. This is consistent with our interpretation of the asymmetric speed of adjustment as evidence of local market power for intermediaries. All of these differential effects are significant at the 99% level. The amount of retail participation slightly lowers the speed of adjustment in both directions, and so does not affect the asymmetry with a high level of statistical significance. Overall, the results in Table IX show that the asymmetric speed of adjustment is associated with variables we would expect to be related to the local market power that broker-dealers exercise.

B. Search Costs and Asymmetric Price Dispersion

The link between asymmetric price adjustment and monopoly power for vendors has been studied theoretically in the context of consumer markets. Classic economic theory predicts that competitive firms immediately pass shocks to their costs through to prices. Asymmetric price adjustment therefore suggests something other than pure competition. Standard models of monopoly or oligopolistic pricing, however, would also generally imply symmetric responses to variation in costs. In this section we consider how models with costly search produce asymmetric price dynamics, and we test the implications these models have for price dispersion in our setting. We find that price dispersion is asymmetric across rising and falling markets, as the models predict. The variation in price dispersion is reversed across the prices at which customers buy from dealers and those at which they sell to dealers, suggesting there are search costs at work on both sides of the market.

In other contexts, search costs have been used to explain price dispersion in OTC markets and hidden costs in financial services. Carlin (2009) describes how opacity in financial markets can be interpreted as intermediaries imposing gratuitous search costs on consumers, and how this sustains monopoly profits in equilibrium. Green (2007) uses search costs to explain price dispersion in settings such as the municipal market studied here, and shows that even when intermediaries must compete for issuers' business, the resulting monopoly rents can be sustained. Recent theories of consumer market behavior show that search costs, or opaque prices, can also lead prices to rise faster than they fall. Thus, our findings of similar behavior here reinforce the importance of search costs in understanding OTC markets.

A number of theories have been advanced to explain the rockets-and-feathers behavior of prices in consumer markets. All of these theories share assumptions that conform to an OTC market, such as that for municipal bonds, where sophisticated institutions trade with less informed retail investors through

broker-dealers. The models all rely on search costs for consumers, and they imply that in equilibrium firms are able to exploit these costs and earn monopoly profits.

For example, Yang and Ye (2008) and Tappata (2009) develop simultaneous search models where search intensity varies with consumers' beliefs about whether firms' costs are high or low. In these models, the prices that firms charge are bounded below by firms' costs (the competitive price) and above by a fixed reservation value for consumers (the monopoly price). If the potential surplus (the distance between producer cost and consumer reservation value) decreases, then equilibrium price dispersion must fall. There is less "room" over which firms can distribute prices. Thus, when costs are low and potential surplus is high, there is more dispersion in prices, and consumers respond by increasing the intensity of search.

The dynamics in price adjustment arise because costs are stochastic but persistent. When costs rise from low to high levels in their models, consumers are already searching intensely, so they quickly learn that prices are rising. When costs drop, few consumers are searching, and, therefore, it takes them several periods to learn about the change and increase search intensity. In equilibrium firms respond by delaying price reductions and spreading prices out. Thus, average prices rise faster than they fall.

Note that by these arguments, price dispersion depends on market conditions as well as average price levels. Our purpose here is to test this implication of the search-based models. One important challenge is coming up with sensible proxies in our setting for the measures of potential surplus that drive the dynamics in the theories. We show below that several variables that can plausibly be interpreted as related to this surplus are related to price dispersion in the way the models predict.

The measures of price dispersion we employ are the absolute dollar dispersion, defined as the daily high minus low price across all transactions at the ask and bid, and the percentage dispersion, defined as the dollar dispersion normalized by the average ask/bid price. Since price dispersion can be observed only when several transactions take place in a given bond on a given day, we perform the analysis both unconditionally and conditional on observing price dispersion. In the latter case, we capture the endogenous nature of trade with a standard Heckman-type sample selection model in which the data are conditioned on days with price dispersion and the sample selection is assumed to follow a probit specification. Appendix C contains details on how we construct the estimator.

Table X shows estimates from a generalized least squares panel regression of our dispersion measures on Treasury yields, changes in Treasury yields, and the Treasury to municipal yield spread. These panel regressions include bond-specific random effects and additional control variables. The table also reports results of the same regression with a Heckman correction for the self-selection associated with observing positive price dispersion and bond-specific fixed effects. Price dispersion for municipal bonds is asymmetric across rising and falling markets. Moreover, the asymmetries are reversed for customer

Table X
Price Dispersion in Rising and Falling Markets

The table documents price dispersion in customer purchases and customer sales under different market conditions. Absolute dollar dispersion is measured as the daily high minus low price across all transactions at the ask or, respectively, bid. Percentage dispersion is dollar dispersion normalized by the average ask/bid price. The column "Pooled" reports coefficient estimates from a generalized least squares panel regression with bond-specific random effects. In this specification, bond-days with and without price dispersion are pooled. The column "Heckman" reports estimates from a panel regression with sample selection in which the data are conditioned on days with price dispersion. The number of observations, therefore, differs across the two specifications. The sample selection is captured by the inverse Mill's ratio from a first-stage probit specification. All specifications include additional controls, for which results are tabulated in Part B of the Internet Appendix. These are the bond and issuer characteristics, order flow variables, and macro variables (term premium, default premium) described in Appendix A, dummies for the calendar year and the U.S. state of issuance, and control variables for the average par size traded. In parentheses we report cluster-robust t -statistics for $H_0 : \beta = 0$. Standard errors are adjusted to account for two-way clustering on bond issuer and calendar time. The estimation sample is the panel of all transactions in the municipal bond market between May 2000 and October 2006, aggregated at the daily frequency.

	Price Dispersion in Customer Purchases				Price Dispersion in Customer Sales			
	%-Dispersion		%-Dispersion		%-Dispersion		%-Dispersion	
	Pooled	Heckman	Pooled	Heckman	Pooled	Heckman	Pooled	Heckman
Treasury yield y_{t-1}^T	Coef. (t) 2.19 (15.04)	Coef. (t) 10.00 (14.92)	Coef. (t) 2.31 (14.77)	Coef. (t) 15.41 (22.66)	Coef. (t) -0.35 (-2.85)	Coef. (t) -14.12 (-8.90)	Coef. (t) -0.16 (-1.32)	Coef. (t) -4.60 (-3.18)
Yield log-spread $\ln(y_t^T / y_t^M)$	-3.03 (-5.04)	-66.40 (-15.23)	-4.95 (-7.92)	-97.77 (-24.18)	1.25 (2.74)	79.56 (9.12)	-0.22 (-0.47)	25.08 (3.10)
Change in Treasury yield Δy^T :								
t	-1.64 (-1.73)	10.63 (4.00)	-1.53 (-1.53)	15.53 (5.69)	-0.92 (-1.60)	-10.71 (-2.62)	-0.63 (-1.09)	-0.31 (-0.08)
$t - 1$	0.81 (0.86)	9.68 (3.87)	0.95 (0.96)	8.20 (3.81)	-0.66 (-1.19)	-5.47 (-1.51)	-0.58 (-1.04)	-4.19 (-1.15)
$t - 2$	1.11 (1.31)	7.35 (3.07)	1.35 (1.51)	8.20 (3.44)	-0.35 (-0.61)	1.10 (0.29)	-0.23 (-0.40)	1.92 (0.50)
Observations	3,592,237	639,626	3,592,237	639,626	3,167,088	246,216	3,167,088	246,216

sales and customer purchases. Consistent with the prediction that investors' search intensity and price dispersion depend on market conditions and proxies for investors' surplus, we find that price dispersion in customer purchases is higher when Treasury yields are high (the absolute price level is low), when the yield spread between Treasuries and munis is low (the surplus is high since the municipal bonds are priced low relative to Treasury bonds), and when interest rates rise (prices fall). We find the exact opposite for the dispersion in customer sales.¹⁵

Dealers in the municipal bond market hedge their inventory positions with Treasuries. Thus, for customers buying bonds, the price of similar Treasuries provides information about dealer costs. When these costs are low (and the Treasury yield is high) we see more dispersion in the prices at which customers buy. Prices at which customers sell are more concentrated when Treasury yields are high, as we would expect if the Treasury price bounds the set of possible prices the seller can expect from the dealer.

We can also regard the yield spread as a measure of potential surplus. When the spread of the Treasury yield over the municipal yield is low, municipal prices are low relative to Treasury prices, suggesting more potential surplus for municipal bond buyers, and less for sellers. As the second row of the table shows, the price dispersion is higher when the potential surplus is higher.

Finally, when interest rates rise and prices fall, we move from a regime with less dispersion and less search to one in which search has greater value for buyers. In the models of search and asymmetric price adjustment, since search intensity is initially low, it takes time for customers to recognize the new regime, leading to greater price dispersion as well as higher average prices. These behaviors for dispersion are evident in the third row of Table X. Again, the fact that the results are reversed for customers buying and selling is quite striking.

V. Conclusion

The consequences of limited transparency in financial markets form a central set of concerns for regulators, investors, and financial intermediaries. We show that opacity in the municipal bond market affects the dynamic behavior of prices. Price discovery is slow, and prices rise faster than they fall. Intermediaries appear to opportunistically time their responses to new information about fundamentals in the prices at which they trade with investors.

We show that dealer markups on inventory positions increase faster when prices rise than they decrease when prices fall. Implicit bid-ask spreads adjust slowly when they are relatively wide and adjust quickly when they are relatively narrow. Implicit half-spreads respond more quickly to price movements

¹⁵Note that we estimate significant positive coefficients on the Inverse Mill's ratio throughout, which suggests that the occurrence of trading and the magnitude of price dispersion are positively related. Sample selection bias is particularly relevant for percentage dispersion in customer sales, as the sign on the yield log-spread switches after bias correction.

when this benefits dealers. Yield spreads relative to Treasuries also adjust with asymmetric speed when they suggest that municipal prices should rise versus fall.

Taken together, these findings suggest that intermediaries benefit from the lack of price transparency and decentralization, and thus from the search costs imposed on investors, in this important OTC market. This implication is reinforced by additional findings that asymmetries are more pronounced for bonds issued in states with segmented markets due to preferential tax treatment for in-state bonds. The link to search frictions for investors is also strengthened by evidence of asymmetry in intraday price dispersion.

Appendix A: Explanatory Variables

The empirical analysis employs the following set of explanatory variables and controls.

- Bond issue and issuer characteristics: issue size; coupon; a zero-coupon dummy; modified duration; maturity in years; indicators for callable, tax-exempt, and insured bonds; dummy variables for each rating category (unrated, AAA, AA, A, BBB, BB-B, CCC-D); and indicator variables for the type of issue and issuer (revenue bond, certificate of participation, tax revenue bond, industrial development bond, housing bond, health care bond, utility bond, facilities bond, tobacco settlement bond, school district issuer, financial authority issuer, development authority issuer).
- Indicator variables for the U.S. state of issuance and calendar year fixed effects.
- Order flow variables: bond-level trading volume and aggregate order imbalances over the past 90 sessions, and state-level and market-wide levels of volume, order imbalances, and new issuance activity over the past 2 weeks.
- Macro variables: short-term interest rate, term premium, and default spread.
- Controls for the average par size traded on a given day, or for the par sizes and the daily changes in the size of the trades used to measure muni yields, or bid, ask, and midpoint prices/yields

Appendix B: Price Dispersion in Rising and Falling Markets

We are interested in how dealers' price setting behavior affects equilibrium price dispersion in the municipal bond market. For this purpose we study how market conditions affect the dispersion in observed prices while controlling for other determinants. One concern is that when there is thin trading, the equilibrium amount of price dispersion may be much larger than the observed distribution simply because observing price dispersion requires trade. For instance, when there are only one or two trades in a given bond on a particular day, we are likely to observe no price dispersion even though dealers may be

quoting vastly different prices to different clients at the same time. On the other hand, the incidence of trade may be correlated with the same variables that determine the amount of price dispersion. For instance, there may be more price dispersion and more or less trading in falling/rising markets or in bonds with certain characteristics. The same may be true for observable determinants, introducing correlation in the error terms. To address this sampling issue we estimate the determinants of price dispersion in two stages. In the first stage, we determine how various market and bond characteristics affect the incidence of observed price dispersion. In the second stage, we estimate how the amount of price dispersion changes in rising versus falling markets and how it depends on various other determinants while conditioning on a sufficient number of trades such that we observe trades occurring at different prices.

Denote by y_{it}^* the (unobserved) equilibrium amount of price dispersion in bond i on day t . Dealers' price setting behavior is such that

$$y_{it}^* = \beta' x_{it} + \epsilon_{it},$$

where x_{it} are observable determinants and ϵ_{it} is an error term with variance σ_ϵ^2 . We are interested in estimating β . We have access only to y_{it} , the price dispersion observed in the sample of trades. Let d_{it} be a binary variable that equals one when y_{it} is strictly positive. The endogenous variable d_{it} depends mostly on the number of trades but also on other characteristics. Let $d_{it} = 1$ whenever $d_{it}^* = \gamma' w_{it} + \eta_{it} > 0$, where w_{it} are observable determinants and η_{it} is an error term of unit variance; that is, correlated with ϵ_{it} . Assume $(\epsilon_{it}, \eta_{it})$ is bivariate normal with correlation ρ . Then by the properties of the normal distribution,

$$E[y_{it}^* | d_{it} = 1] = \beta' x_{it} + E[\epsilon_{it} | \eta_{it} > -\gamma' w_{it}] = \beta' x_{it} + \rho \sigma_\epsilon \frac{\phi(\gamma' w_{it})}{\Phi(\gamma' w_{it})}.$$

The model can thus be estimated by first estimating the coefficients γ in a probit regression and then estimating the parameters of interest, β , by regressing y_{it} on x_{it} and on the inverse Mill's ratio $\lambda_{it} = \frac{\phi(\gamma' w_{it})}{\Phi(\gamma' w_{it})}$. The latter term corrects for the bias induced by the data sampling, and the coefficient on the inverse Mill's ratio measures the correlation ρ between the error terms, scaled by σ_ϵ .

REFERENCES

- Balduzzi, Pierluigi, Edwin J. Elton, and T. Clifton Green, 2001, Economic news and bond prices: Evidence from the U.S. Treasury market, *Journal of Financial and Quantitative Analysis* 36, 523–543.
- Bessembinder, Hendrik, William F. Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251–288.
- Carlin, Bruce I., 2009, Strategic price complexity in retail financial markets, *Journal of Financial Economics* 91, 278–287.
- Ederington, Louis H., and Jae Ha Lee, 1993, How markets process information: News releases and volatility, *Journal of Finance* 48, 1161–1191.

- Ederington, Louis H., and Jae Ha Lee, 1995, The short-run dynamics of the price adjustment to new information, *Journal of Financial and Quantitative Analysis* 30, 117–134.
- Edwards, Amy K., Lawrence Harris, and Michael S. Piwowar, 2007, Corporate bond market transparency and transactions costs, *Journal of Finance* 62, 1421–1451.
- Fleming, Michael J., and Eli M. Remolona, 1999, Price formation and liquidity in the U.S. Treasury market: The response to public information, *Journal of Finance* 54, 1901–1915.
- Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri, 2007, Transparency and liquidity: A controlled experiment on corporate bonds, *Review of Financial Studies* 20, 235–273.
- Green, Richard C., 2007, Issuers, underwriter syndicates, and aftermarket transparency, *Journal of Finance* 62, 1529–1550.
- Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007a, Financial intermediation and costs of trading in an opaque market, *Review of Financial Studies* 20, 275–314.
- Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007b, Dealer intermediation and price behavior in the aftermarket for new bond issues, *Journal of Financial Economics* 86, 643–682.
- Harris, Lawrence, and Michael S. Piwowar, 2006, Secondary trading costs in the municipal bond market, *Journal of Finance* 61, 1361–1397.
- Kidwell, David S., Timothy W. Koch, and Duane R. Stock, 1984, The impact of state income taxes on municipal borrowing costs, *National Tax Journal* 37, 551–562.
- Li, Dan, 2007, Price clustering and dealer's market power in OTC markets, Working paper, Tepper School of Business, Carnegie Mellon University.
- Neumark, David, and Steven A. Sharpe, 1992, Market structure and the nature of price rigidity: Evidence from the market for consumer deposits, *Quarterly Journal of Economics* 107, 657–680.
- Peltzman, Sam, 2000, Prices rise faster than they fall, *Journal of Political Economy* 108, 466–502.
- Piazzesi, Monika, 2005, Bond yields and the Federal Reserve, *Journal of Political Economy* 113, 311–344.
- Schultz, Paul, 2001, Corporate bond trading costs: A peek behind the curtain, *Journal of Finance* 56, 677–698.
- Tappata, Mariano, 2008, Rockets and feathers: Understanding asymmetric pricing, *RAND Journal of Economics* 40, 673–687.
- Yang, Huanxing, and Lixin Ye, 2008, Search with learning: Understanding asymmetric price adjustments, *RAND Journal of Economics* 39, 547–564.