The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXVIII, NO. 2 • APRIL 2023

Do Municipal Bond Dealers Give Their Customers "Fair and Reasonable" Pricing?

JOHN M. GRIFFIN, NICHOLAS HIRSCHEY, and SAMUEL KRUGER[*]

ABSTRACT

Municipal bonds exhibit considerable retail pricing variation, even for same-size trades of the same bond on the same day, and even from the same dealer. Markups vary widely across dealers. Trading strongly clusters on eighth price increments, and clustered trades exhibit higher markups. Yields are often lowered to just above salient numbers. Machine learning estimates exploiting the richness of the data show that dealers that use strategic pricing have systematically higher markups. Recent Municipal Securities Rulemaking Board rules have had only a limited impact on

[*]John Griffin and Samuel Kruger are with the McCombs School of Business, University of Texas at Austin. Nicholas Hirschey is with the Nova School of Business and Economics, Universidade NOVA de Lisboa. We thank Scott Bauguess; Hank Bessembinder; Thummim Cho; Daniel Garrett; Patrick Herb; Edith Hotchkiss; Stacey Jacobson; Stefan Nagel (Editor); Dermot Murphy; Zhaogang Song; Norman Schürhoff; Chester Spatt; Kumar Venkataraman; Stephane Verani; two anonymous referees; an Associate Editor; and seminar participants at London Business School, HKUST, Hong Kong Polytechnic University, the University of Texas, the Nova School of Business and Economics, the University of Melbourne, the 2019 Lone Star Finance Conference, the 2019 Northern Finance Association meeting, the 2019 Adam Smith Juniors Workshop, the 2020 American Finance Association meeting, the 2020 European Winter Finance Conference, the 2020 Midwest Finance Association meeting, and the 2021 Municipal Finance Workshop. We also thank Brady Wyllie for research support, and we are grateful for research support from London Business School, Integra Research Group, a University of Texas McCombs Research Excellence Grant, Nova School of Business and Economics, Fundação para a Ciência e a Tecnologia (UIDB/00124/2020, UIDP/00124/2020, and Social Sciences DataLab - PINFRA/22209/2016), POR Lisboa, and POR Norte (Social Sciences DataLab, PINFRA/22209/2016). Griffin is an owner of Integra Research Group, Integra FEC, and Integra REC, which engage in research and financial consulting on a variety of issues related to investigating fraud, including various types of bonds. Disclosure statement is available on *The Journal of Finance* website. Supplementary results can be found in an Internet Appendix, which is available in the online version of this article.

Correspondence: Nicholas Hirschey, Nova School of Business and Economics, Universidade NOVA de Lisboa, Rua da Holanda 1, 2775-405 Carcavelos, Portugal; e-mail: nicholas.hirschey@novasbe.pt.

DOI: 10.1111/jofi.13214

markups. While a subset of dealers focus on best execution, many dealers appear focused on opportunistic pricing.

FINANCIAL FIRMS FREQUENTLY STATE A COMMITMENT to low pricing and prioritizing their clients' best interests. For example, Vanguard has a stated commitment to low fees and "clients first," Merrill Lynch advertises "putting your investing needs, wants and priorities first," and Edwards Jones advertises "zero competing interests…clients come first." Along the same lines, in 2019 the Securities and Exchange Commission (SEC) adopted Regulation Best Interest, which codifies a responsibility to pursue retail customers' best interests. Yet not all pricing and fees are readily transparent, and financial services firms face a conflict of interest because a commitment to best pricing and low fees can come at the expense of firm profits and employee bonuses. Thus, in practice, do financial firms put customer interests first? The answer to this question is unclear in part because areas of finance with opaque and discretionary pricing are not fully understood.

We study the pricing practices of financial firms in a large area of finance with potential conflicts of interest and limited transparency for retail investors—the municipal bond market. This market consists of over $3.9 trillion in outstanding debt and over $1.4 trillion of secondary market trading per year (Wu and Vieira (2019)). Unlike many other active markets, customer trading is decentralized among dealer networks with electronic trading platforms only in the interdealer market. In contrast to other over-the-counter securities, municipal bonds are frequently sold to retail investors because of their attractive tax features. Whereas households own only 6% of outstanding corporate bonds, they own 51% of municipal bonds (Bessembinder, Spatt, and Venkataraman (2020)). Though regulations require "fair and reasonable" pricing, brokers directly profit from markups. This conflict may be mitigated by regulation and oversight from the Municipal Securities Rulemaking Board (MSRB), which has the explicit goal of ensuring fair and efficient pricing.

Our analysis delivers four primary findings. First, municipal bond markups to retail customers have remained high and variable over the past 15 years despite significant regulatory efforts to enhance transparency and improve execution quality since high markups were first noted in the municipal bond literature. Second, municipal bond purchase prices frequently vary significantly even for the same bond sold on the same day, and even by the same dealer within the day, which casts doubt on compliance with regulatory requirements for fair and consistent pricing. Third, dealers use practices that may exploit investors' limited attention and cognitive biases such as using round prices and yields, targeting yields to stay above salient thresholds, and differentially marking up long-maturity bonds, all of which are associated with higher markups. Fourth, markups vary widely across dealers, and dealers' past propensity to engage in strategic pricing practices and their market share are strong predictors of individual bond markups.

What drives markups and markup dispersion? In their comprehensive survey on bond market trading, Bessembinder, Spatt, and Venkataraman (2020) conclude by noting that one of the most fundamental unanswered questions is why bond trading costs are relatively high.[1] The literature proposes several nonmutually exclusive explanations for high municipal bond markups, including lack of transparency (Harris and Piwowar (2006)), costs accumulating as multiple dealers intermediate a transaction (Schultz (2012)), dealers with cost advantages due to network centrality charging higher customer markups (Li and Schürhoff (2019)), and use of market/bargaining power with customers (Green (2007), Green, Hollifield, and Schürhoff (2007b)).

We start by comparing small, retail-sized municipal bond markups over time from 2005 through 2019. In the new issue market, median markups and extreme markups have not decreased since real-time trade reporting was established in 2005. Median markups in the seasoned issue market have decreased slowly over time.[2] Yet, like the new issue market, the seasoned issue market still exhibits considerable extreme markups of 3% or more, even in 2019.

Using MSRB's proprietary data set with dealer IDs for municipal bond trades from July 2011 through December 2017, we find that some dealers consistently charge close to 0% markups over their cost on small trades in both the new issue and the seasoned issue market, but 53% (47%) of dealers charge median markups of more than 1% in the new (seasoned) issue market. MSRB rules G-18 and G-30 require that transactions traded on the same day should generally receive consistent pricing. Yet when comparing retail-sized purchase transactions in the same bond on the same day, we find substantial variation in prices and dealer markups. In the new issue market, the difference between the 10% most expensive and 10% least expensive purchase prices for small trades is at least 0.5% on 61% of trading days. Dispersion of at least 1% occurs 44% of the time. Price differences are similar in the seasoned issue market. Dealer fixed effects for transactions of the same bond on the same day differ by as much as 2 percentage points (ppt), even for bonds with extremely active trading. Dealers also vary pricing practices across small customers. While a few dealers employ similar markups on all of their trades, many charge widely varying markups. Even when considering the same dealer selling the same bond on the same day, purchase price differences between the 10% most expensive and 10% least expensive small trades of at least 0.5% occur 35% of the time and differences of at least 1% occur 18% of the time. This result holds even when restricting the sample to transactions of the exact same size and hour of the day, and the practice is more prevalent for dealers with higher markups.

To shed light on dealer motivations and practices driving markups, we investigate the specific prices and yields of the municipal bond trades. Municipal bond investors care about yields, and institutional municipal bond trades

[1] Biais and Green (2019) find that municipal bond trading costs on the NYSE in 1926 to 1927 were around half of recent costs.

[2] This result is consistent with the findings of Chalmers, Liu, and Wang (2021) and Wu and Vieira (2019) that secondary market trading costs have decreased significantly over time, but these studies do not examine new issue pricing.

are executed mainly at exact basis-point yields. However, trades can also be executed based on price. We find that when municipal bonds trade based on yields at basis-point increments, they exhibit markups that are considerably lower. Trading frequently clusters at quarter and odd-eighth percentage-point yield increments, and trades at these yield increments tend to have higher markups, indicating that the practice of rounding is associated with high markups. Clustering at price increments of quarters and odd-eighths of a dollar is also associated with higher markups. These clustering patterns are most pronounced for small trades, but they are also present to a lesser degree in medium and large trades, particularly at odd-eighth price increments. These patterns are highly significant, vary across bonds, and are strongly related to dealer fixed effects. Christie and Schultz (1994) find that NASDAQ market makers collusively avoided trading on odd-eighths to increase spreads over 25 years ago. Though not necessarily collusive, quoting municipal bonds at coarse price and yield increments indicates that dealer pricing choices contain a substantial discretionary component and may be a mechanism to keep spreads and markups high. Consistent with this interpretation, markups and price clustering are not lower when more dealers trade a bond.

To further disentangle pricing practices, we look to other features of the data that may reflect strategic pricing. A growing literature shows that firms may profit by setting prices to exploit customers' limited attention, cognitive biases, or lack of financial sophistication (Ellison (2006) and Gabaix and Laibson (2006)). For example, firms frequently set prices to exploit the left-digit effect, whereby customers perceive numbers such as 1.01 to be significantly higher than 0.99.[3] Consistent with this prediction, we find strong discontinuities in trade frequencies around exact percentage-point yields. For example, yields of 3.01% are much more common than yields of 2.99%. Moreover, these trades have higher markups than trades with yields just below exact percentage points. This result suggests that dealers use their price discretion to raise markups and decrease yields without crossing salient thresholds that customers would notice. This pricing practice is much more prevalent for small trades and among high-cost dealers.

Another form of strategic pricing could be charging higher markups on longer-maturity bonds because markups have a smaller and less noticeable effect on these bonds' yields. For matched pairs of bonds within the same issue, long-maturity bonds have 1.25 ppt higher markups than their short-maturity bonds counterpart, even after requiring the long-maturity bonds to be more liquid. This practice is again consistent with charging higher markups when they are less salient to customers.

What is driving cross-sectional differences in markups? We predict markups based on dealer and bond characteristics using a gradient-boosting decision tree machine learning model that exploits the richness of the trade-level data

[3] The left-digit bias is a long-standing result in marketing (Thomas and Morwitz (2005)) and has also been documented in accounting (Carslaw (1988)), bank deposit yields (Kahn, Pennacchi, and Sopranzetti (1999)), and stock trading (Bhattacharya, Holden, and Jacobsen (2012)).

and allows for potential nonlinearities and interactions between bond and dealer characteristics. Gradient-boosting decision trees and neural networks tend to be the best machine learning methods for tabular data, and decision trees have the additional advantage that they require less parameter tuning and are relatively easy to interpret (Chollet (2021), Howard and Gugger (2020)). Dealer-level rounding and market share variables can predict markups almost as well as knowing the dealer's average markup over the past 30 days. The model predicts higher markups for bonds purchased from dealers who engage in pricing practices that exploit behavioral biases or who have large market share. We also find that dealers who were likely part of the bond's syndicate at issuance typically charge higher seasoned issue markups.

We find no evidence that median or extreme markups decreased following MSRB rule changes requiring best execution in March 2016, even for high-cost dealers. We find some small decreases in seasoned issue markups following regulatory changes that increased markup transparency in May 2018, but no decreases in the new issue market. Duflo (2017) emphasizes the importance of program details in effective program design. Practical challenges to accessing pricing information for typical retail investors, ambiguities in MSRB regulatory language, and few (12 over five years) enforcement actions may help explain the lack of regulatory effectiveness. Over the 6.5-year sample, the total value of small trade markups (for purchase transactions) in excess of 2% is \$479 million, with \$1.83 billion of markups in excess of 1% and \$2.83 billion of markups in excess of 0.5%, indicating that markups are economically sizeable.

Our findings demonstrate widely varying pricing practices across dealers. Some dealers deliver low markups and consistent pricing to their customers, but many charge high and variable markups. Many dealers seem more concerned with maximizing markups through customer cognitive biases and limited attention than with providing best execution. Given the negligible economic impact of MSRB rule changes, significant municipal bond market reforms, such as proposals by Harris (2015) and Harris, Kyle, and Sirri (2015) for customer access to electronic trading, may be needed to improve pricing.[4] A growing literature documents conflicts of interest and misconduct among brokers and financial advisors (e.g., Egan (2019), Egan, Matvos, and Seru (2019), Dimmock, Gerken, and Graham (2018)). Consistent with this literature, our evidence suggests that customers should approach brokers with a high level of caution.

The next section describes the municipal bond market and the regulatory environment. Section II describes the data and summary statistics. Section III assesses pricing consistency. Section IV assesses differences across dealers using regression analysis. Section V identifies granular pricing practices

[4] Hendershott and Madhavan (2015) find that access to electronic auctions decreases trading costs in the corporate bond market, and Hau et al. (2021) find that retail foreign exchange over-the-counter clients with access to platform trades receive drastically lower transactions costs.

used by dealers to increase markups. Section VI estimates machine learning models to predict markups and pricing practices. Section VII discusses the impact of MSRB rule changes and other potential explanations. Finally, Section VIII concludes.

I. Municipal Bond Market and Regulatory Environment

A. Market Background

The municipal bond market is an important market to study due to its size, customer base, and regulatory features. Municipal bond interest is exempt from federal income taxes and is typically exempt from state income taxes in the issuer's state. This creates a natural clientele consisting of individuals with high marginal tax rates in the state of the bond's issuance (Bergstresser and Cohen (2015), Babina et al. (2021)). When a new bond is issued, its underwriter (or underwriting syndicate) sells the bond to either its own customers or to other dealers, who in turn sell the bond to other dealers or their customers.[5] As discussed by Schultz (2012), selection of a particular bond is almost always broker initiated, not customer initiated. Municipal bonds are frequently held for long holding periods or until maturity. Thus, trading volume is highest for newly issued bonds.

Harris and Piwowar (2006) and Green, Hollifield, and Schürhoff (2007a) find that small municipal bond trades have substantially higher markups than large trades. Green, Hollifield, and Schürhoff (2007b) use a structural model and find that the high trading costs in municipal bonds likely come from dealer market power. The time period of their study (2000 to January 2004) preceded substantial MSRB reforms, and prices were reported only at the end of the day, and even less frequently in some cases. Green, Li, and Schürhoff (2010) find that prices rise faster than they fall, which is also consistent with market power. Harris and Piwowar (2006) indicate that the largest reason for high prices is a lack of transparency.

Transaction price transparency has increased over time. Since January 2005, all municipal bond transactions have been publicly posted with a 15-minute lag. Nonetheless, it is not clear how aware less sophisticated customers are of this resource, particularly given Schultz's (2012) finding that this increased transparency has had little impact on average markups. By contrast, increasing transparency with the Trade Reporting and Compliance Engine (TRACE) system had a large effect on corporate bond transaction costs (Bessembinder, Maxwell, and Venkataraman (2006), Edwards, Harris, and Piwowar (2007), Goldstein, Hotchkiss, and Sirri (2007), and Bessembinder and Maxwell (2008)). To facilitate access to post-trade transaction information, the MSRB launched the Electronic Municipal Market (EMMA) website in March

[5] Recent studies of municipal bond underwriting include Bergstresser and Herb (2022) and Garrett (2021).

2008 and has improved the general functionality of the website over time (Wu and Vieira (2019)).[6]

B. MSRB Regulatory Framework

The municipal bond market is under the jurisdiction of the SEC with much left to self-regulation by the MSRB. The MSRB's goal is to ensure that the municipal bond market is "fair and efficient." An important component of fair pricing is the MSRB's best execution rule (MSRB Rule G-18), which requires that dealers obtain the most favorable possible pricing for customers given current market conditions. Similarly, MSRB Rule G-30 requires that commissions and markups for transactions with dealers must be "fair and reasonable." The MSRB's rulemaking guidance and the SEC's enforcement activities have clarified that "a 'fair and reasonable' price bears a reasonable relationship to the prevailing market price of the security," which is best determined by close comparisons to other transactions (MSRB Rule G-30, Supplementary Material, .02(a) and (c)). Guidance for both rules emphasizes that trades on the same day should generally have the same price.[7]

A recent rule change that took place on May 14, 2018 requires dealers "to disclose mark-ups and mark-downs … to retail customers on certain principal transactions and to provide dealers guidance on prevailing market price for the purpose of determining mark-ups" (MSRB Regulatory Notice 2016-28). The MSRB's stated goal is to enhance compliance with the MSRB's fair-pricing rules by providing retail customers and dealers with additional useful price information.

C. Possible Explanations for Variation in Markups

The institutional features of the municipal bond market give rise to several potential reasons for markup variation. First, markups may vary due to differences in execution cost structure. Different dealers may acquire the same bond at different prices. Accordingly, we focus on markups over dealer cost in most of our analysis and include various trade size controls. Second, markups could vary with bond and trade characteristics due to factors such as liquidity. Analyzing the new issue market, where bonds are more actively traded, mitigates this concern. In addition, we control for bond characteristics that might

[6] The literature on corporate bonds also finds elevated markups for small trades (Schultz (2001), Edwards, Harris, and Piwowar (2007)). Goldstein, Hotchkiss, and Nikolova (2021) document significant corporate bond price dispersion for small trades. Harris (2015) finds that corporate bond transactions often occur at prices that are inferior to available quotes.

[7] "A transaction chain that results in a large difference between the price received by one customer and the price paid by another customer for the same block of securities on the same day, without market information or news accounting for the price volatility, raises the question as to whether each of these customers received a price reasonably related to the market value of the security" (MSRB Rule G-30, Supplementary Material, .04(a)).

be associated with liquidity, and we find that pricing varies within and across dealers even for trades of the same bond on the same day.

Alternatively, markups may vary due to dealer pricing practices. Dealers may vary in their commitment to fair and consistent pricing, and customers could have different levels of sophistication and knowledge about bond prices. While it is difficult to ascertain dealer motives, dealers who focus on giving clients fair pricing and higher yields may quote bonds in terms of yields instead of prices. By contrast, dealers who focus on profit margins or have clients that do not pay as much attention to yields may focus instead on the price they charge investors and thus could be more likely to quote bonds in prices. Dealers may also vary in the extent to which their trades with small customers cluster at round prices and yields. Clustering at round prices is common in other asset markets, including stocks, in part because round numbers operate as salient anchors due to cognitive biases (Shiller (2000), Bhattacharya, Holden, and Jacobsen (2012)). Round prices facilitate easier price negotiation (Harris (1991)), but they also make pricing coarser, which increases bid-ask spreads (Christie and Schultz (1994)). Li (2007) finds interesting evidence of price, yield, and markup clustering with higher markups for rounded trades in a municipal bond sample from April 2002 to January 2004, before many market reforms, which she interprets as consistent with dealers exercising market power.[8]

Dealers who focus on maximizing their markups may try to increase prices (thereby lowering yields) when investors are less likely to notice. A broker who might otherwise quote a bond at 3.09% or even 3.18% might adjust the price up, moving the yield down and close to but not below 3.00%. This behavior would result in more trade activity at yields of 3.01% or 3.02% than at 2.99% or 2.98%. Another potential framing that could affect markup saliency is related to the fact that markups may be less noticeable on bonds with long maturities because markups have less impact on stated yields as bond maturity increases. Hence, brokers may charge higher markups on longer-maturity bonds. We describe our exact tests in more detail in the empirical sections below.

II. Data, Summary Statistics, and Preliminaries

A. Sample Construction

Our primary data source is MSRB academic data from July 2011 through December 2017. This data set includes buying and selling dealer fields for each trade record that are populated with anonymized dealer IDs. In keeping with the literature, we limit the sample to customer purchases. A trade's markup is a standard measure of transaction costs. With the detailed dealer data, we calculate the markup as $\frac{\text{Customer Purchase Price}}{\text{Dealer's Purchase Cost}} - 1$. Matches between a customer's trade and the dealer's average purchase price are made following the first in, first out (FIFO) procedures of Green, Hollifield, and Schürhoff (2007b) applied

[8] Similarly, Goldstein, Moser, and Van Ness (2022) find price clustering in corporate bonds, particularly for small trades.

to each dealer's inventory. To reduce noise from stale matches, the seasoned issue sample is limited to matches with no more than seven days between matched purchases and sales. Overall, there are 2.8 million small trades in the new issue market and 8.2 million small trades in the seasoned issue market (as shown in Table IA.IV along with additional summary statistics).[9] To put markups in a longer-term perspective, we also calculate markups from January 2005 to December 2019 in the expanded Wharton Research Data Services (WRDS) data set, which has a longer time series but no dealer IDs. In this sample, the new issue markup is the percent premium over the offering price. Additional details on the data and filters are provided in the Internet Appendix.

B. Markups over Time

Over the past 15 years, the municipal bond market has changed considerably, including increased transparency, more awareness of the EMMA website, increased accessibility, growing public attention to markups, government fines and regulatory changes related to excessive markups, and a shift toward electronic trading in the interdealer market (Cestau et al. (2018)). Panel A of Figure 1 plots new issue markups over offering prices and seasoned issue round-trip markups for small, medium, and large trades from January 2005 to December 2019. Following Green, Hollifield, and Schürhoff (2007b), small trades are defined as trades with a par value of \$100,000 or less, medium trades have a par value between \$100,000 and \$500,000, and large trades have a par greater than \$500,000. Markups in the new issue market have no clear trend. For small trades, the median markup starts just below 1% in 2005, rises to around 2% in the early 2010s, and then falls to just below 1% in 2019. Median new issue markups for medium and large trades are 0% throughout 2005 to 2019.

To gauge the incidence of high markups, we also plot the 95$^{\text{th}}$ percentile of markups over time. We do not find much decline in high new issue markups, which seems surprising. The 95$^{\text{th}}$ percentile markup for small new issue trades is over 3% throughout the sample and was about the same in 2019 as in 2005. The 95$^{\text{th}}$ percentile markup for larger new issue trades also did not change much between 2005 and 2019. It was mostly between 2% and 3% for medium trades and between 1% and 2% for large trades.[10]

In the seasoned issue market, median markups for small trades were typically above 1.5% before 2008, rose to around 2% from 2008 to 2014, and have fallen since 2014, consistent with Chalmers, Liu, and Wang (2021), Wu and Vieira (2019), and Bessembinder, Spatt, and Venkataraman (2020). However,

[9] The Internet Appendix may be found in the online version of this article.

[10] These findings extend Schultz's (2012) finding that new issue markups did not change much from 1999 to June 2010, following the introduction of trade reporting in January 2005. The findings are also consistent with Bessembinder et al.'s (2018) analysis of corporate bonds, which shows relatively flat trading costs from 2006 to 2016.

Figure 1. Markups over time and across underwriters. Panel A plots quarterly 50th and 95th percentiles of trade markups. Panel B shows 50th and 95th percentiles of new issue markups on small ($\leq$\$100k) customer purchases for deals by top underwriters. Underwriters are sorted left to right from low to high markups. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. Trade sizes are determined by the trade's total par amount: small is $\leq$\$100k, medium is $>$\$100k and $\leq$\$500k, and large is $>$\$500k. The new issue markup is the percentage premium over the offering price. The seasoned issue markup is the estimated total spread captured by the dealer sector when matching customer sales to eventual customer purchases. The markup is (Customer Purchase Price/Matched Dealer Cost) – 1. The seasoned issue matches consist of the combined set of immediate, round-trip, and FIFO matches calculated following Green, Hollifield, and Schürhoff (2007b). We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. For the underwriter sample, we include primary market ("List Offering/Takedown") trades and secondary market trades in the first 14 days of trading. We include the top 30 underwriters by small-trade dollar volume in 2019. There are more than 30 underwriter names because when there is a merger, we include the premerger version of the 2019 underwriter as a separate observation (e.g., "BA-Corp" and "BAML"). (Color figure can be viewed at wileyonlinelibrary.com)

the 95th percentile of markups is still high. The 95th percentile was around 4% from 2005 to 2008, spiked to over 5% in 2009, and then decreased to around 3% by 2019 for an overall decrease of only around 1 ppt since 2005. Similar to the new issue market, medium- and large-trade markups are much lower.

C. Markups across Underwriters

Underwriters play an important role in the new issue market and may influence markups. Panel B of Figure 1 summarizes how new issue markups vary across underwriters by plotting the median markup categorized by the underwriter of the bond. Because this analysis does not require dealer IDs, we use the WRDS data set to obtain a longer time series and to include takedown primary market trades (which are excluded from the dealer ID data set). Markups vary substantially across underwriters.[11] Bonds underwritten by Fidelity, Roosevelt Cross, Bernardi Securities, Baird, Jeffries, and Siebert all have median markups of approximately zero. In contrast, bonds underwritten by Mesirow, U.S. Piper, Boenning, and PNC Capital have median small-trade markups near 2%, with several other large underwriters just below them. Controlling for bond and trade characteristics shows similar differences across underwriters (Figure IA.4). The substantial variation across bond underwriters suggests that underwriters play an important role in influencing markups either through their pricing strategies or through their choices of which dealers to work with.

D. How Much Do Markups Vary across Trade Sizes?

The time trends in Figure 1 clearly show large pricing differences between small, medium, and large trade sizes. In this section, we assess differences in markups across trade sizes by regressing markups and bond prices on indicators for trade size ranges after controlling for observed and unobserved differences in the bonds being traded. Table I reports results for new issues in Panel A and seasoned issues in Panel B. In column (1) of Panel A, we regress new issue markups on indicators for small and medium trade sizes with control variables for bond characteristics as well as state, month, credit rating, and days-since-offering fixed effects. Standard errors are three-way clustered by dealer, bond, and day.[12] Large trades are the omitted category. The most important control variable is maturity, with a coefficient of 45 basis points (bps) for log maturity. After controlling for bond characteristics, small trades have markups that are 75 bps higher than large trades, and medium trades are also

[11] Figure IA.3 summarizes the lead underwriters that are most active in small trades. Baird is the largest underwriter, with over $1B of small trade volume in 2019, followed by Stifel and Bank of America.

[12] We use three-way clustering for standard errors throughout the paper whenever sufficient data are available. Table IA.VI considers alternative standard error clustering and shows that three-way clustering is the most conservative.

Table I

Markup Regressions

This table examines the difference between markups of small (≤\$100k), medium (>\$100k and ≤\$500k), and large (>\$500k) trades. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) – 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. Both new issue and seasoned issue regression control variables include ln(bond's maturity), ln(bond's total par issued), and dummy variables indicating the bond's credit rating on a scale from 0 (unrated) to 24 (AAA), the bond's state, the month of the trade, insured bonds, general obligation bonds, callable bonds, bonds with a sinking fund, bank qualified bonds, taxable bonds, and AMT bonds. In addition, new issue bonds have a dummy for negotiated offerings and dummies for each day since the bond was offered to the public, while seasoned issues add a control for ln(seasoning in days since issuance). All variables are winsorized at the 0.5% and 99.5% levels. Three-way clustered (dealer, bond, and day) standard errors are in parentheses. Significance levels are indicated by asterisks for p-values less than or equal to 0.01 (***), 0.05 (**), and 0.1 (*).

Panel A: New Issues							
	Markup (%)				ln(Price) × 100		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
isSmall	0.748***	0.511***	0.380***	0.126***	0.505***	0.372***	0.128***
	(0.088)	(0.057)	(0.054)	(0.043)	(0.056)	(0.052)	(0.043)
isMedium	0.390***	0.312***	0.239***	0.063**	0.302***	0.224***	0.063**
	(0.051)	(0.035)	(0.039)	(0.030)	(0.034)	(0.037)	(0.030)
LnMaturity	0.445***						
	(0.048)						
LnSizeBond	0.015						
	(0.016)						
Insured	0.086**						
	(0.036)						
GeneralObligation	−0.032*						
	(0.018)						
Callable	0.315***						
	(0.060)						
SinkingFund	0.112**						
	(0.056)						
BankQualified	−0.078**						
	(0.035)						
Taxable	0.128***						
	(0.035)						
AMT	0.137***						
	(0.052)						
Negotiated	0.130***						
	(0.035)						
State, Month, and Rating FE	✓						
Days Since Offering FE	✓						
Bond-Day FE		✓	✓		✓	✓	
Dealer FE			✓			✓	
Dealer-Bond-Day FE				✓			✓
Observations	3,428,687	3,428,687	3,428,687	3,428,687	3,428,687	3,428,687	3,428,687
Adjusted R^2	0.407	0.751	0.789	0.910	0.993	0.994	0.997
Dependent Variable Mean	1.21	1.21	1.21	1.21	463.6	463.6	463.6

Continued

Table I
Continued

Panel B: Seasoned Issues							
	Markup (%)				ln(Price) × 100		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
isSmall	0.777***	0.471***	0.345***	0.173***	0.423***	0.306***	0.165***
	(0.070)	(0.045)	(0.035)	(0.032)	(0.041)	(0.031)	(0.032)
isMedium	0.385***	0.316***	0.219***	0.101***	0.302***	0.209***	0.097***
	(0.036)	(0.034)	(0.027)	(0.025)	(0.033)	(0.025)	(0.025)
LnMaturity	0.473***						
	(0.041)						
LnSizeBond	−0.029***						
	(0.010)						
Insured	0.040**						
	(0.016)						
GeneralObligation	−0.028*						
	(0.015)						
Callable	0.122***						
	(0.029)						
SinkingFund	0.170***						
	(0.045)						
BankQualified	0.012						
	(0.021)						
Taxable	0.087*						
	(0.051)						
AMT	0.148***						
	(0.051)						
LnSeasoning	−0.110***						
	(0.012)						
State, Month, and Rating FE	✓						
Bond-Day FE		✓	✓		✓	✓	
Dealer FE			✓			✓	
Dealer-Bond-Day FE				✓			✓
Observations	9,351,768	9,351,768	9,351,768	9,351,768	9,351,768	9,351,768	9,351,768
Adjusted R^2	0.355	0.786	0.826	0.902	0.996	0.996	0.998
Dependent Variable Mean	1.17	1.17	1.17	1.17	467.1	467.1	467.1

marked up 39 bps relative to large trades. In column (2), we control for potential unobserved differences across bonds and over time by adding bond-day fixed effects. This specification conceptually compares trades in the same bond on the same day, with the result that small trades have markups that are on average 51 bps higher compared to large trades.

We also consider specifications that compare trade sizes sold by the same dealer. In column (3), we add dealer fixed effects, and in column (4) we add dealer-bond-day fixed effects to compare trades of the same bond on the same day sold by the same dealer. In both specifications, small trades again have significantly higher markups.

Markups are a function of both the price of the bond and the dealer's cost to acquire it. To examine whether differences in markups are due to pricing rather than acquisition cost differences, columns (5), (6), and (7) repeat the same regressions with log price as the dependent variable instead of markups.

Consistent with the markup regressions, small transaction prices are 0.51% higher with bond-day fixed effects, 0.37% higher with bond-day and dealer fixed effects, and 0.13% higher with dealer-bond-day fixed effects.[13]

Panel B of Table I repeats the analysis in the seasoned issue market with almost identical results. The overall takeaway is that small trades routinely have higher markups.

III. How Much Does Pricing Vary across and Within Dealers on the Same Day?

In this section, we assess whether pricing is consistent across small trades. Observed and unobserved bond characteristics may vary across bonds and over time, so we begin our examination by shutting down these channels, that is, by considering purchases of the same bond on the same day. As we discuss above, MSRB rules G-18 and G-30 require that transactions traded on the same day generally receive consistent pricing. To assess same-day pricing discrepancies, we restrict the sample to small purchases of the same bond on the same day.

A. Same-Bond, Same-Day Trade Comparisons

We require five small-trade observations per bond on the same day, resulting in 1.8 million new issue trades and 1.3 million seasoned issue trades. We first examine trades in the same bond on the same day. We next turn to trades in the same bond on the same day from the same dealer. Finally, we consider trades that also have the exact same trade size. The analysis examines both price and markup dispersion.

The left side of Figure 2 plots the distribution across bond-day observations of differences between the average price of the 10% most expensive and 10% least expensive small purchases of the same bond on the same day.[14] The bars are a histogram of price dispersion, and the line is the complementary cumulative distribution function, which shows the percentage of observations that have dispersion above a particular level. If customers receive identical pricing on the same bond bought on the same day, these markup differences should be zero. Instead, prices frequently vary within same-bond, same-day groups of transactions. The difference between the 10% most expensive and 10% least expensive prices is at least 0.5% for new issues 61% of the time. Even more extreme dispersion of at least 1% occurs 44% of the time. Price differences are similar in the seasoned issue market.

Do dealers treat all of their small customers the same, or is there heterogeneity in markups that dealers charge their customers? The middle column

[13] In Table IA.VII, we adjust acquisition costs for changes in the S&P National AMT-Free Municipal Bond Index while the bond is in a dealer's inventory, with almost identical results.

[14] For bond-days with between 5 and 10 trades, this is the difference between the maximum and minimum price.

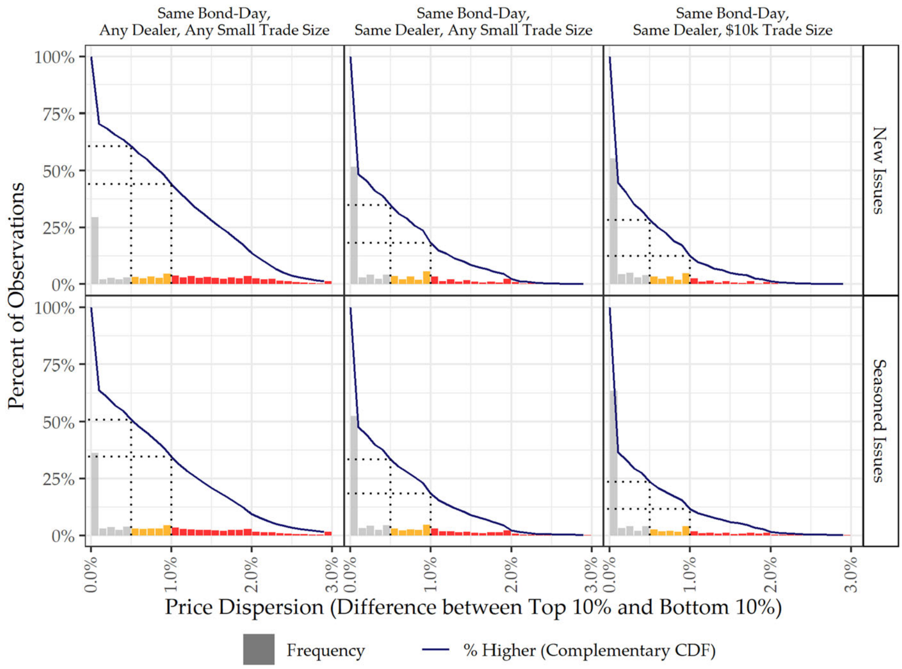

Figure 2. Same bond, day, dealer, and trade size: Price dispersion. This figure illustrates price differences for trades on the same bond-day, same dealer-bond-day, or same dealer-bond-day that also have the same trade size. The plots show the frequency of price differences between the 10% most expensive and 10% least expensive trades for trades from the same bond-day(-dealer)(-trade size) grouping. In each plot, the bars show the frequency of price differences and the line shows the percent of those observations for which the difference is greater than a given threshold (i.e., the complementary cumulative distribution function). We limit the sample to small trades and require at least five trades per grouping. When there are fewer than 10 trades in a grouping, the difference is between the most and least expensive trades. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

of Figure 2 plots price dispersion for the same bond on the same day with the additional restriction that all transactions are with the same dealer. The unit of observation is groups of five or more small purchases of the same bond on the same day from the same dealer. If dealers give customers identical pricing on small purchases of the same bond bought on the same day, these price differences should be zero. We find, however, that a difference between the 10% most expensive and 10% least expensive prices for new issue bonds from the same dealer of at least 0.5% occurs 35% of the time, and a difference of at least 1% occurs 18% of the time.[15] Price differences are similar in the seasoned issue market.

Is this within-dealer markup dispersion due to different trade sizes? In the right plots of Figure 2, we further restrict the analysis to trades with a trade size of exactly $10,000, which is the most common small-trade size. Once again, we find that significant dispersion of at least 0.5 ppt occurs 28% of the time in new issues and 24% of the time in seasoned issues.

Is price dispersion for small trades greater than for medium and large trades? We compare dispersion of residual standard deviations across trade size groups from regressions with markup or price as the dependent variable. Across new and seasoned issue markets in specifications with bond-day, dealer-bond-day, or dealer-bond-day-trade-size fixed effects, price and markup dispersion are consistently higher for small trades than for medium or large trades (Figure IA.7).

We also consider several other possible explanations. First, we examine whether price differences are potentially driven by dealers acquiring the bond at different costs by repeating Figure 2 using markup differences instead of price differences. The results, which are shown in Figure IA.8, are virtually identical. Next, we examine whether the within-day price differences are caused by changes in the bonds' market values over the course of the day. Comparing trades within the same hour instead of the same day leads to almost identical results (Figure IA.9). In summary, markup differences represent different prices for the same security from the same dealer at essentially the same time, which would seem to be a clear failure of pricing fairness according to MSRB regulations and guidance.

IV. Regression Analysis

To estimate dealer fixed effects in a regression framework, we restrict the sample to small trades. In addition to dealer and bond-day fixed effects, we also control for log trade size. The regression is equivalent to column (3) of Table I with the restricted sample and additional trade size control variable. The dealer fixed effects capture dispersion in dealer markups relative to other dealers selling the same bond on the same day.[16] If different dealers are giving customers trading the same bond similar pricing, dealer fixed effects should be close to zero.

The top plot of Figure 3 shows that in the new issue market, fixed effects

[15] When examining within-dealer markup differences more generally without restricting attention to trades in the same bond on the same day, the median dealer has an almost 2 ppt difference between its 90$^{\text{th}}$ and 10$^{\text{th}}$ percentile markup, and markup differences persist even after controlling for bond and trade characteristics (Figures IA.5 and IA.6).

[16] To center the fixed effects around zero, we also impose the restriction that all of the fixed effects must sum to zero. Standard errors are three-way clustered by dealer, bond, and day with a Bonferroni correction for multiple comparisons. The Bonferroni correction adjusts the t-statistic thresholds for the number of fixed effects being estimated. For dealer fixed effects, a p-value of 0.05 corresponds to a t-statistic of 3.60.

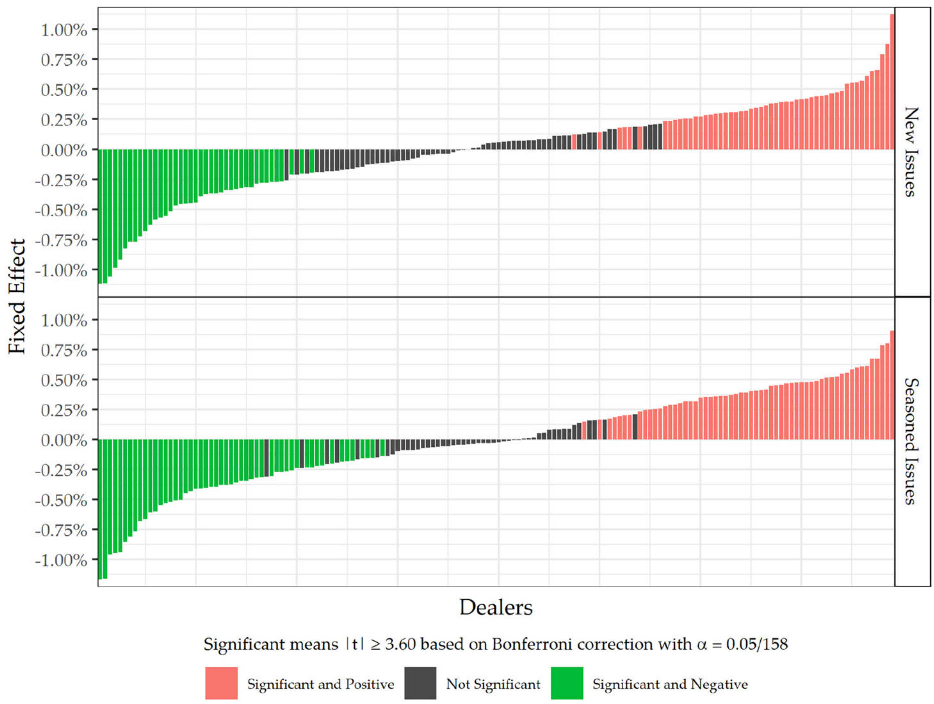

Figure 3. Dealer markup adjusting for bond-day fixed effects. This figure shows dealer fixed effects from regressions in which the dependent variable is a trade's markup and the sample is restricted to small trades. We display fixed effects for the top dealers, that is, dealers who have at least \$20 million worth of trades with customers in both the new and seasoned issue markets where the trade size was ≤\$100k. We calculate dealer fixed effects using the following regression: $New\ Issue\ Markup_i = \alpha + \Sigma_{k=1}^{N_{dealers}} \beta_k Dealer_{i,k} + \beta_1 LnParTraded_i + \Sigma_{b=1}^{N_{bonds}} \Sigma_{d=1}^{N_{days}} \beta_{b,d} BondDay_{i,b,d} + \epsilon_i$,$Seasoned\ Issue\ Markup_i = \alpha + \Sigma_{k=1}^{N_{dealers}} \beta_k Dealer_{i,k} + \beta_1 LnParTraded_i + \Sigma_{b=1}^{N_{bonds}} \Sigma_{d=1}^{N_{days}} \beta_{b,d} BondDay_{i,b,d} + \epsilon_i$, where *Dealer* and *BondDay* are dummy variables. We subtract the mean top dealer fixed effect from all others to center the distribution on zero and plot the result. Standard errors are three-way clustered by dealer, bond, and day. The null hypothesis is that the demeaned fixed effect equals zero. Statistical significance is determined using the Bonferroni correction such that a fixed effect is significant if the p-value is less than 0.05# Dealers. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost)−1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

range widely from less than negative 1% to more than 1%, and 92 of the 158 brokers have fixed effects that are significantly different from zero at the 5% level. The bottom plot shows a similar pattern for the seasoned issue market. Both markets exhibit a small group of dealers on the far left of the plot with large negative fixed effects. These are dealers who consistently deliver lower markups. There are also several dealers at the far right, with fixed effects greater than 0.5%. Differences across dealers are similar when we examine price differences rather than markup differences and when the sample is restricted to highly liquid bonds based on the bonds' size and trading volumes (Figure IA.10 and Table IA.VIII). For additional context on markup magnitudes, we also plot dealers' median raw markups, which show that many dealers have median markups of less than 25 bps while other dealers have median markups above 2% (Figure IA.11).[17] Dealers with high markups also have more pricing variation across trades of the same bond on the same day (Figure IA.15 and exhibit larger markup differences between their small and large trades (Figure IA.16). Overall, the findings indicate that there are considerable and persistent differences across dealers in their markups even for the same bond trading on the same day.

V. Why Do Markups Vary?

Same-bond, same-day variation in bond pricing cannot be explained by differences in bond risk or other bond characteristics. We next turn to granular data on the price points and yields of individual trades to assess pricing practices related to trading at round increments on price or yield. We then turn to strategic pricing around salient yield thresholds and to bond maturity.

A. *Is Markup Variation Related to Price and Yield Pricing Patterns?*

Price points and yields of individual trades could relate to dealer motivations and trading practices. The granularity of the MSRB data, which record yields with 10^{th} of a basis point precision and prices with 10^{th} of a penny precision, allows us to differentiate trades quoted in yields from trades quoted in prices by looking at the precise price and yield of the trade. Yields at exact basis-point increments indicate that a trade was likely quoted in yields, and prices at exact penny increments indicate that a trade was likely quoted in prices. In the new issue market, trades at the offering price are dropped from the analysis because their pricing is driven by the offering price as opposed to a dealer decision.

[17] To assess whether these markup differences are due to dealers purchasing bonds at different prices, the top plot in Figure IA.11 calculates new issue markups relative to offering price and shows dispersion is similar. To assess whether these differences between dealers are more consistent with random variation or a persistent dealer trait, we examine the correlation between dealers' median markups in the first half and the second half of the sample and find correlations of 0.696 in new issues and 0.706 in seasoned issues (Figure IA.13). Dealers specializing in retail trades (i.e., dealers with more small-trade volume relative to their large-trade volume) tend to have higher markups across trade sizes, but this explains only a small amount of variation across dealers (Figure IA.14).

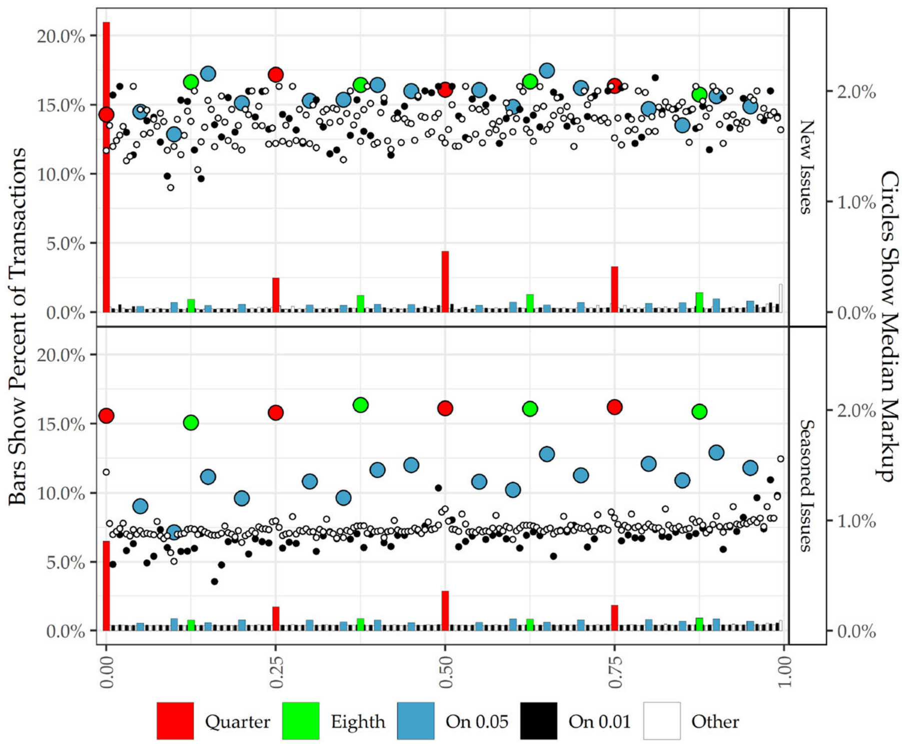

Figure 4. Price rounding: Frequency and median markups. This figure plots a histogram of small (≤$100k) trade frequencies and markups for trades binned by fractional dollar prices (e.g., $100.035 → 0.035). Price bins use half-penny increments for the histogram. We plot markups for trades that lie on exact increments with colored dots. Markups for trades that do not lie on exact increments are white. New issue trades at the offering price are excluded. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

A.1. Pricing Patterns

Figure 4 shows a histogram of small-trade prices by half-penny (0.005) bin sizes with bin frequencies plotted as bars on the left axis and median markups for each bin plotted as circles on the right axis. The plotted prices are modulus after subtracting off dollars. New issues are in the top plot, and seasoned issues are in the bottom plot. Red bars and circles represent prices on the quarters, odd one-eighths are in green, and trades on other five-cent increments are in blue. Trades on penny price increments are in black and those not on any of these increments are in white.

Several patterns emerge. First, even though the data specify prices to one-tenth of a penny, trades in the new and seasoned issue markets frequently cluster at quarters. There is more clustering at dollar and 50-cent increments, particularly in the new issue market. Clustering on odd-eighths and every five cents also occurs, but not as frequently. Second, trades that cluster at quarter, odd-eighth, and, to a lesser extent, five-cent increments have higher median markups. Third, the difference in markups is large, especially in seasoned issues, where trades on pennies and nonspecial "Other" increments have median markups around 1%, whereas trades on the quarter increments and odd-eighths increments hover a full percentage-point higher at 2% markups.

What might drive the clustering? Clustering around quarters and odd-eighths even in the seasoned issue market, where prices are not typically related to issuance price, indicates that prices likely contain a discretionary component. Coarser increments may be a mechanism chosen by dealers to facilitate higher spreads and markups. Competition does not alleviate price clustering; on average, prices are actually somewhat coarser when competition is high (Table IA.XX). Panel A of Table II shows the prevalence of price clustering in small trades compared to medium and large trades. Coarse prices are defined as quarter and odd-eighth prices, and fine prices are defined as nonquarter exact-penny prices. In the new (seasoned) issue market, small trades occur at coarse prices 32.4% (13.5%) of the time, compared to 18.6% (7.4%) for medium and 8.7% (5.7%) for large trades. While price clustering is less common for medium and large trades, it is still associated with elevated markups (Figure IA.18).

A.2. Yield Patterns

Because bond investors mainly care about yields, municipal bonds are often quoted and thought of in terms of yields as opposed to price. Figure 5 plots histograms and markups for small-trade yields similar to the price plots discussed above. The new issue plot again excludes trades at the offering price. Several interesting patterns emerge. First, markups vary significantly with yields. Most strikingly, markups are significantly lower when yields are at exact basis-point increments (plotted in black) compared to yields that are not at the basis point or other round increments (plotted as small white dots). Second, there is substantial clustering on yields at the quarters, odd-eighths, and five-basis-point increments. Third, compared to basis-point yields (the black dots), trades at yields on the quarters and odd-eighths have significantly elevated markups.[18] Interestingly, bonds trading at odd-eighths have higher markups than bonds trading on the quarters, particularly in the seasoned issue market.

Panel B of Table II summarizes yield clustering for small, medium, and large trades. Coarse yields are defined as quarter and odd-eighth percentage-point

[18] Figure IA.17 aggregates and summarizes these markup differences.

Table II

Trade Rounding Frequencies

This table examines the frequency of price and yield rounding for transactions according to trade size. The rounding statistics (Coarse Price/Yield and Fine Yield) are based on transactions that are not at the offering price. Transactions at coarse prices and yields are defined as those at exact quarters and odd-eighths. Transactions at fine yields are at nonquarter exact basis-point yields. To do so, we focus on fractional dollar prices (e.g., \$100.035 → 0.035) and fractional percentage yields (e.g., 4.512% → 0.512).

	Small Trades [0, \$100k]	Medium Trades (\$100k, \$500k]	Large Trades > \$500k
	Panel A: Prices		
New Issues			
% Coarse	32.4	18.6	8.7
% Fine	13.1	12.4	11.3
% Other	54.5	69.0	80.0
Seasoned Issues			
% Coarse	13.5	7.4	5.7
% Fine	15.9	12.8	11.3
% Other	70.6	79.8	83.0
	Panel B: Yields		
New Issues			
% Coarse	13.5	13.6	9.1
% Fine	18.4	46.5	69.7
% Other	68.1	39.9	21.2
Seasoned Issues			
% Coarse	9.6	12.2	9.3
% Fine	24.5	48.5	65.6
% Other	65.9	39.2	25.1

yields, and fine yields are defined as nonquarter exact basis-point yields. The most striking feature of the data is that large new (seasoned) issue trades occur at fine yields 69.7% (65.6%) of the time, compared to only 18.4% (24.5%) of the time for small trades. This result likely reflects the use of yields for more sophisticated customers. Medium- and large-customer trades are also more likely to cluster at quarter and five-basis-point increments, and medium and large trades at odd-eighth yields again have some of the highest markups in both the medium and the large size groups (Figure IA.19).

A.3. Price and Yield Clustering Regressions

We now turn to regressions of markups on price and yield clustering to jointly examine different types of clustering and to control for bond and trade characteristics. Results for small trades are reported in Table III with new issue trades in columns (1) to (4) and seasoned issue trades in columns (5) to (8). Standard errors are three-way clustered by dealer, bond, and day. Focusing first on new issue trades, column (1) reports results from regressing markups on indicators for price and yield increments with no control variables.

Table III

Rounding Regressions

This table examines the relationship between small ($\leq$\$100k) trade markups and dummy variables indicating price and trade rounding. New issue regressions exclude trades at the offering price. The dependent variable is a trade's markup in percent. Both new issue and seasoned issue regression control variables include ln(trade's par amount), ln(bond's maturity), ln(bond's total par issued), and dummy variables indicating the bond's credit rating on a scale from 0 (unrated) to 24 (AAA), the bond's state, the month of the trade, insured bonds, general obligation bonds, callable bonds, bonds with a sinking fund, bank qualified bonds, taxable bonds, and AMT bonds. In addition, new issue bonds have a dummy for negotiated offerings and dummies for each day since the bond was offered to the public, while seasoned issues add a control for ln(seasoning in days since issuance). We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. Three-way clustered (dealer, bond, and day) standard errors are in parentheses. Significance levels are indicated by asterisks for p-values less than or equal to 0.01 (***), 0.05 (**), and 0.1 (*).

	New Issues				Seasoned Issues			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Variables								
PriceAt01OrEighth	0.111***	0.045***	0.041**	0.048***	−0.046	−0.058	−0.055	0.029**
	(0.017)	(0.017)	(0.018)	(0.019)	(0.062)	(0.039)	(0.036)	(0.014)
PriceAtOff05	0.045	−0.063	−0.181***	−0.206***	0.265***	0.088	0.053	−0.080
	(0.108)	(0.091)	(0.067)	(0.065)	(0.088)	(0.068)	(0.056)	(0.053)
PriceAtOddEighth	0.187*	0.083	0.003	−0.001	0.670***	0.362***	0.263***	0.066
	(0.096)	(0.085)	(0.073)	(0.078)	(0.130)	(0.101)	(0.085)	(0.076)
PriceAtQuarter	0.147*	0.019	−0.070	−0.098	0.703***	0.295***	0.197**	−0.052
	(0.084)	(0.062)	(0.056)	(0.062)	(0.138)	(0.106)	(0.088)	(0.069)

Continued

Table III
Continued

	New Issues				Seasoned Issues			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yield Variables								
YieldAt01OrEighth	−0.617***	−0.341***	−0.154***	−0.080***	−0.423***	−0.200***	−0.150***	−0.112***
	(0.154)	(0.067)	(0.041)	(0.027)	(0.093)	(0.064)	(0.052)	(0.037)
YieldAtOff05	−0.113	−0.056	−0.049	−0.015	0.181**	0.156***	0.126***	0.079
	(0.153)	(0.096)	(0.038)	(0.038)	(0.086)	(0.056)	(0.045)	(0.052)
YieldAtOddEighth	0.538***	0.320***	0.099**	0.024	0.617***	0.315***	0.216***	0.139***
	(0.176)	(0.082)	(0.048)	(0.034)	(0.085)	(0.044)	(0.029)	(0.035)
YieldAtQuarter	0.276*	0.140**	0.037	−0.014	0.369***	0.226***	0.177***	0.120***
	(0.161)	(0.069)	(0.047)	(0.036)	(0.108)	(0.072)	(0.056)	(0.043)
Controls		✓	✓	✓		✓	✓	✓
Days Since Offering FE		✓	✓					
Exclude Trades at Offering Price	✓	✓	✓	✓				
State FE		✓				✓		
Month FE		✓				✓	✓	
Rating FE		✓				✓	✓	
Bond FE			✓				✓	
Bond × Day FE				✓				✓
$\geq$ 5 Bond × Day Trades				✓				✓
Observations	2,350,562	2,350,562	2,350,562	1,508,829	8,225,408	8,225,408	8,225,408	1,305,120
Adjusted R^2	0.101	0.321	0.570	0.621	0.073	0.351	0.466	0.774
Dependent Variable Mean	1.63	1.63	1.63	1.81	1.25	1.25	1.25	1.80

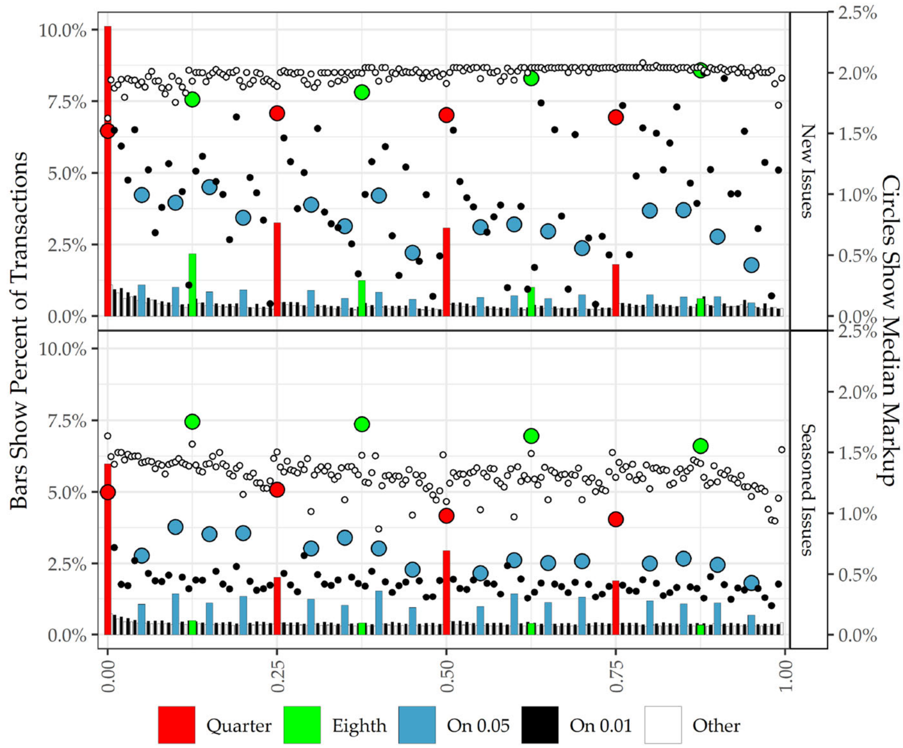

Figure 5. Yield rounding: Frequency and median markups. This figure plots a histogram of small (≤$100k) trade frequencies and markups for trades binned by fractional percentage yields (e.g., 4.512% → 0.512). Yield bins use half-basis-point increments for the histogram. We plot markups for trades that lie on exact increments with colored dots. Markups for trades that do not lie on exact increments are white. New issue trades at the offering price are excluded. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

The results are largely consistent with Figures 4 and 5, indicating that examining price and yield simultaneously does not substantively affect inferences. Trades at penny or eighth price increments have higher markups compared to nonpenny trades, and trades at the odd-eighths and quarters have additionally higher markups (significant at the 10% level). Exact basis-point yields are associated with lower markups. Compared to trades at nonround exact basis-point yields, odd-eighth percent yield trades have elevated markups. Trades at quarter percent yields also have elevated markups, but the coefficient is statistically significant at only the 10% level.

To examine whether differences in quoting practices are explained by the types of bonds being traded, column (2) includes bond and trade characteristic controls variables as well as state and month fixed effects. In the new issue market, bonds traded on penny or eighth price increments have slightly higher markups, but prices at the odd-eighths and quarters have no incremental effect, likely indicating that the effects of coarser, more clustered pricing on markups is related to bonds of certain characteristics. We next examine differences in markup quoting within a particular bond by including bond fixed effects in column (3). It is natural that the relation between markups and clustering should be dampened because competition will push dealers to quote in a similar fashion to other dealers within a bond. Nevertheless, clustering is still related to within-bond markup variation. Quoting on penny or eighth price increments is associated with an extra 4.1 bps markup, and quoting on basis-point or eighth yields is associated with a decrease in markups of 15.4 bps. Column (4) focuses on trading variation within a bond during the same trading day by including bond-day fixed effects. Here, we restrict the sample to those days that have at least five transactions. The magnitudes of pricing differences are substantially reduced in this specification because prices for the same bond on the same day are easily compared to one another. Nonetheless, there are still statistically significant effects indicating that bonds traded on penny or eighth prices have higher markups while bonds trading on basis point or eighth yields have lower markups.

Columns (5) to (8) report results from the same regression specifications in the seasoned issue market. Unlike in the new issue market, penny prices have little relation to markups in the seasoned issue market. Trades with odd-eighths prices and quarter prices have elevated markups except in the bond-day fixed effect specifications. Basis-point yields are associated with lower markups in all specifications. Relative to basis-point yield trades, trades at odd-eighth and quarter percent yields have elevated markups across all specifications. A similar pattern holds for 0.05 ppt yield increments across all regressions except the bond-day fixed effects specification. Overall, both the new issue and the seasoned issue market show consistent relationships between markups and price and yield clustering.

A.4. Dealer Price and Yield Practices

One possibility is that these patterns are related to dealer practices. Figure 6 plots the fraction of each dealer's trades that cluster at coarse (quarter or odd-eighth) prices or yields on the x-axis against dealer median markups on the y-axis. The color of the dots shows the fraction of each dealer's transactions that are at fine yields, ranging from less than 5% in dark red to over 90% in white. The size of the dots corresponds to three dealer-size groups.

The scatterplot shows wide differences in price and yield clustering across dealers, with some dealers almost never using coarse prices or yields and other dealers pricing at coarse prices or yields nearly all the time. Dealers who use coarse pricing tend to have higher markups, as indicated by the

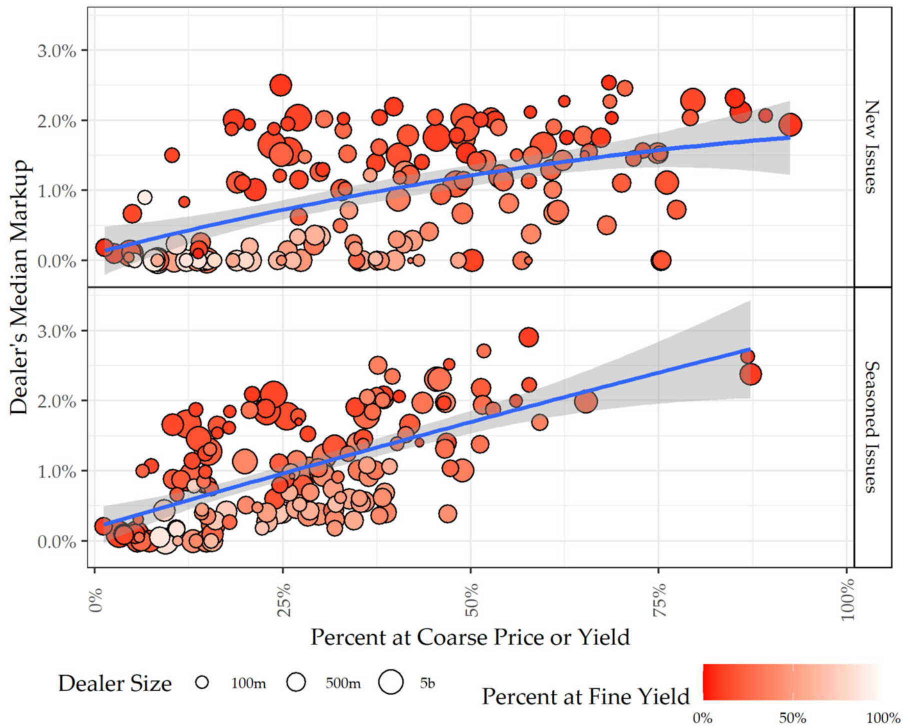

Figure 6. Relation between dealer markups and how often they round prices or yields. This figure shows how a dealer's median markup on small (≤\$100k) trades varies according to the percentage of their transactions occurring at coarse price or yield increments. To do so, we look at the fractional dollar prices (e.g., \$100.035 → 0.035) and fractional percentage yields (e.g., 4.512% → 0.512). On the x-axis, we count the percentage of transactions at either coarse prices or coarse yields. The color of the circle indicates the percentage of trades that are at fine yields. Transactions at coarse prices and yields are defined as those at exact quarters and odd-eighths. Transactions at fine yields are at nonquarter exact basis-point yields. The size of the circle is the total par amount the dealer trades in the new and seasoned issue small trades. In each plot, the blue line and gray shaded area show the OLS fit and 95% confidence interval for markups as a function of percent of transactions at coarse prices or yields. The regressions are fit using a second-order polynomial. New issue trades at the offering price are excluded. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

upward-sloping predicted markup curve based on a regression of dealer median markups on a second-order polynomial of dealer price and yield coarseness. Dealers with fine yields (indicated by lighter dots) tend to have lower markups even among dealers with the same level of pricing coarseness.[19]

[19] In Figure IA.20, we further analyze price and yield rounding for low- and high-cost dealers by replicating Figures 4 and 5 separately for dealers in the bottom and top terciles of small-trade

A.5. Discussion of Price and Yield Clustering

The preceding results support three conclusions. First, customers receive better pricing when trading at exact basis-point yield increments. Trading on basis-point yields is more common for large trades than for retail-sized trades, and low-cost dealers are more likely to trade on basis-point yields. Second, clustering on round prices and yields is common and frequently associated with higher markups. By pricing bonds at round numbers, dealers increase the coarseness of the market, which widens bid-ask spreads. As a result, markups are higher at round prices and yields. Third, differences across dealers suggest that dealer decisions drive the practices above.

B. Is Markup Variation Related to Strategic Pricing?

Are dealers strategically using these pricing practices to maximize markups? Figure 7 plots trade frequency by half-basis-point (0.005) yield increments from 10 bps below to 10 bps above exact percentage-point yields for small trades.[20] There is considerable clustering in the number of trades just above the threshold. In the new issue market, there are over three times as many trades with yields ending in the $(0.005\%, 0.015\%]$ range compared to yields ending in the $(0.985\%, 0.995\%]$ range.[21] In the seasoned issue market, there is also a significant, though smaller, discontinuity at percentage-point yields. These discontinuities are a special feature of small trades. Figure IA.21 replicates the analysis for medium and large trades and shows no discontinuity for large trades and a much smaller discontinuity for medium trades.

Median markups are higher in the $(0.005\%, 0.015\%]$ range compared to the $(0.985\%, 0.995\%]$ range in the seasoned issue market. The difference is 24 bps for nonrounded trades (white circles) and 30 bps for rounded trades (black circles). In the new issue market, markups around the whole-yield threshold are more consistent, but they are still higher above the threshold than below. Median markups are 12 bps higher for nonrounded trades above the threshold and 33 bps higher for rounded trades above the threshold. The higher markups above percentage-point yield thresholds indicate that the mass of trades above

markups. Coarse prices and yields are more prevalent for high-markup dealers. Though clustering is less common for low-markup dealers, their markups are also elevated at coarse prices and yields.

[20] The nearest bars on either side of 0.00 show the mass of trades in the half basis point (0.005) immediately above and below the exact percentage-point yield. The extra mass at exact percentage-point yields (the 0.00 point) is excluded to more easily see the area on either side of the yield-clustering threshold.

[21] We excluded yields ending in $(0.995\%, 0.005\%]$ from this comparison because they will be equal to the even percentage-point yield when rounded to the nearest basis point. The discontinuity can also be seen in the plotted probability density function (along with a 95% confidence interval) implied by regressing trade frequency on a second- order polynomial of yields and dummy variables for each 5 bp increment, again excluding the half-basis-point (0.005) increments immediately above and below the percentage yields.

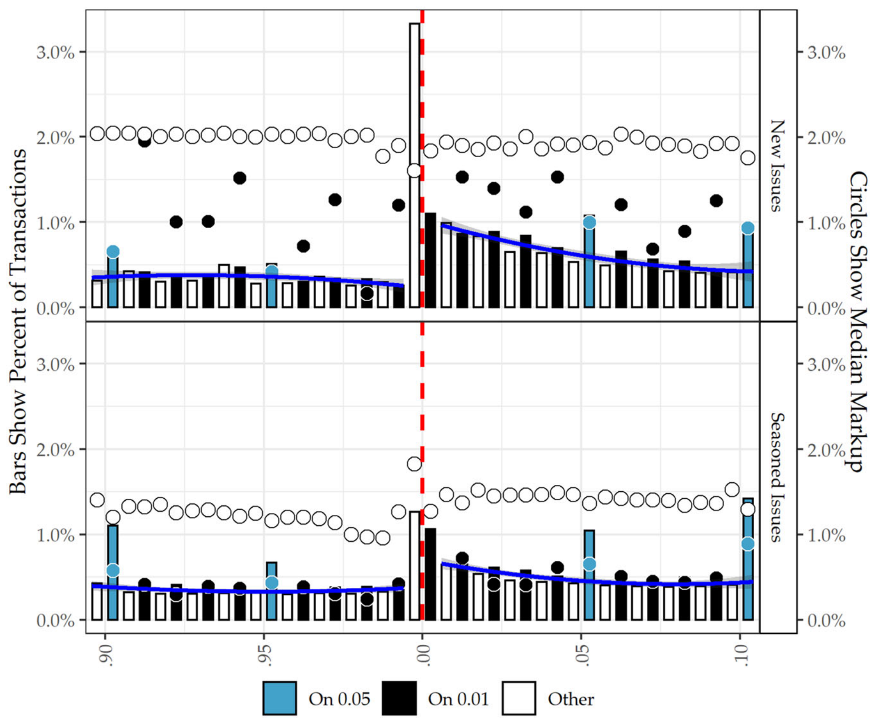

Figure 7. Transactions with yields near a whole percent. This figure shows small (≤$100k) trade frequency by half-basis-point (0.005) increments from 10 bps below to 10 bps above exact percentage-point yields. The nearest bars on either side of 0.00 show the mass of trades in the half basis point immediately above and below the exact percentage-point yield. The extra mass at exact percentage-point yields (the 0.00 point) is excluded to more easily see the area on either side of the yield clustering threshold. Yields are assigned to bins by fractional percentage yields (e.g., 4.512% → 0.512). Colored circles show markups for trades on exact 0.01 and 0.05 increments. White circles show markups for trades that are not priced at these increments. In each plot, the blue lines show the probability density function estimated by regressing trade frequency on a second-order polynomial of yields and dummy variables for each 5 bp increment. The estimated density function is plotted along with a 95% confidence interval. The estimated density does not include the half-basis-point increments immediately above and below the yield threshold because they round to the percentage-point yield. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. (Color figure can be viewed at wileyonlinelibrary.com)

the threshold is not primarily due to dealers shifting prices down to help clients achieve larger yields above the threshold. Rather, dealers are raising prices as much as they can without reducing yields below percentage-point thresholds. Because the resulting yields are above salient left-digit thresholds, dealers may expect the markups to be less noticed by investors, consistent with a large marketing literature.[22]

Another interesting feature of the figure can be seen in the bar exactly below the percentage-point threshold. Because yields are rounded to the nearest basis point for many reporting conventions, yields between 0.995 and 1.00 will round up to 1.00. Thus, these bonds appear as if they have higher percentage-point yields. If dealers are engaged in strategic yield pricing, one should expect these bonds to exhibit similar patterns to those above the threshold. Consistent with this prediction, we find considerable clustering just below the threshold, with markups higher on these trades in the seasoned issue market. By contrast, there is little clustering below the threshold for large trades and much smaller clustering for medium trades (Figure IA.21).

We also examine these patterns separately for low- and high-cost dealers.[23] Consistent with high-cost dealers focusing more on strategic pricing, the patterns are much more pronounced for high-cost dealers in both the new issue and the seasoned issue markets (Figure IA.23).

C. Does Long Maturity Facilitate Higher Markups?

Large markups may be harder for investors to detect in longer-maturity bonds because the effect of markup on yield is decreasing in bond maturity. MSRB interpretive guidance also prioritizes looking at a bond's yield to assess whether its price is fair.[24] Results reported earlier in Table I show that bond maturity is one of the largest drivers of cross-sectional differences in markups. To assess the joint relation that maturity has with markups and pricing coarseness, we plot average markups and pricing coarseness across bond maturities for small trades. As maturity increases, pricing shifts from fine yields to coarse prices, and markups increase even conditional on the type of pricing (Figure IA.24 and Table IA.IX).[25]

To further control for potential liquidity differences, we examine a matched sample of short-maturity (10 years or less) and long-maturity (20 years or more) bonds within the same issue.[26] The average markup difference between the matched long and short-maturity bonds is 1.25 ppt and highly statistically

[22] Figure IA.22 shows results around the quarter percentage-point yields that are not on the percentage point. Patterns are directionally similar but not as strong, as one might expect because percentage-point thresholds are more salient.

[23] Low-cost dealers are the third of dealers with the lowest median markups, while high-cost dealers are the third with the highest median markup, from Figure IA.11.

[24] "The MSRB firmly believes that the resulting yield to the customer is the most important factor in determining the fairness and reasonableness of a price in any given transaction" (MSRB Interpretive Letter for Rules G-21, G-30, and G-32, December 11, 2001).

[25] These effects are also present but much smaller for medium and large trades (Figure IA.25 and Table IA.X).

[26] To ensure the longer-maturity bonds are not less liquid, the longer-maturity bond is required to have at least as much small-trade dollar volume as the shorter-maturity bond in the first 14 days of trading. Further details on the matching procedure are in the caption of Table IV. The matching results in 12,236 matched pairs across 7,433 issues, with a reduced sample for rounding

Table IV

Matched Pairs of Long- and Short-Maturity Bonds in the Same Issue

This table compares new issue small (≤\$100k) trade markups for matched pairs of short- and long-maturity bonds from the same issue. We form pairs of bonds from the same issuance as identified by the same six-digit CUSIP and same issue date. We then match a bond with a maturity of 10 years or less to a bond from the same issue with a maturity of 20 years or more. To make sure the long-maturity bonds are not less liquid, we compare the total par traded in the new issue market. We calculate total par traded only using trades of that trade size. We require the longer-maturity bond to have at least as much par traded as the shorter-maturity bond. The bond size difference is also required to be within five million or 25% of their average size. For each bond with a maturity below 10 years, we match it with the shortest-maturity bond with a maturity over 20 years. If there is another short-maturity bond in the issue, then we match it with the next-youngest bond with a long maturity. We average first across trades in the short- and long-maturity bonds in a pair separately. The table reports averages and standard errors for the mean across all short- and long-maturity bonds in all pairs. Markup is in percent, maturity is in years, and total par traded is in dollars. The rounding statistics (Coarse Price/Yield and Fine Yield) are based on transactions that are not at the offering price; pairs are excluded from rounding calculations if all trades are at the offering price. Transactions at coarse prices and yields are defined as those at exact quarters and odd-eighths. Transactions at fine yields are at nonquarter exact basis-point yields. To do so, we focus on fractional dollar prices (e.g., \$100.035 → 0.035) and fractional percentage yields (e.g., 4.512% → 0.512). A bond's dollar volume is total par traded in small trades over the entire new issue period. Markups are calculated relative to the dealer's cost estimated on a FIFO basis. Two-way clustered (underwriter and month) standard errors are in parentheses. Significance levels are indicated by asterisks for p-values less than or equal to 0.01 (***), 0.05 (**), and 0.1 (*).

Variable	Long Avg.	Short Avg.	Difference	# of Pairs	# of Issues
Markup (%)	1.45	0.21	1.25***	12,236	7,433
			(0.04)		
% at Coarse Price	30.78	5.03	25.75***	7,081	4,738
			(0.68)		
% at Coarse Yield	12.23	11.07	1.15	7,081	4,738
			(0.72)		
% at Fine Yield	16.32	42.88	−26.55***	7,081	4,738
			(1.04)		
Maturity	23.19	4.32	18.87***	12,236	7,433
			(0.11)		
Bond\$Volume	754,129	111,917	642,211***	12,236	7,433
			(28,602)		

significant as shown in Table IV. The long-maturity bonds are also 25.75 ppt more likely to have a coarse (quarter or odd-eighth) price and 26.55 ppt less likely to have a fine (exact basis point, nonquarter percent) yield.[27]

Overall, the evidence suggests that use of different pricing conventions, exploitation of cognitive biases, and mitigated impact on yields due to matu-

results due to dropping transactions at the offering price. Standard errors are two-way clustered by underwriter and month.

[27] We also examine medium and large trade sizes and find that long-maturity bonds have markups that are 0.88 ppt higher for medium trades and 0.20 ppt higher for large trades (Table IA.XI), compared to the difference of 1.25 ppt for small trades.

rity play important roles in small trade markups. The common theme across practices is that certain dealers appear to use their pricing discretion to charge higher markups when investors are less likely to notice.

VI. How Important are Dealer Practices in Explaining Markups?

We now seek to systematically quantify the extent to which dealer characteristics and practices can explain markup dispersion across dealers. For this analysis, we use a flexible machine learning model that estimates decision trees with gradient boosting. Machine learning is well suited to this analysis because of the magnitude of the data and the many potential nonlinearities and interactions involved in predicting and explaining markups.[28] We first define our variables of interest and then introduce and estimate the machine learning model.

A. Dealer Variables

There are many potential explanations for why certain dealers charge high markups. We examine how well measures of a dealer's inventory, market share, and pricing practices explain markups. Inventory could be important because when a dealer's inventory is high, they have an incentive to decrease markups to generate sales that reduce inventory. Market share may be important because when a dealer has high market share, they may be able to charge higher markups because it is costly for customers to search for an alternate counterparty. Strategic pricing practices have a direct effect on markups, and they might also be helpful for identifying dealers engaged in other unobserved practices that increase markups. Specifically, even if a particular trade is not executed at a coarse price, knowing that the selling dealer often executes trades at coarse prices might suggest that the dealer engages in other practices that increase markups. We construct dealer-level variables to measure these effects as described in the Internet Appendix. The variables are calculated separately for new and seasoned issues on a rolling basis using a dealer's trades from the prior 30 days. The *DealerInventory* measure captures a dealer's abnormal inventory relative to the past 30 days aggregated across positions acquired in all bonds and every trade size. The rest of the dealer-level variables are calculated using only small purchases.

B. Gradient Boosting Decision Trees

In addition to dealer and bond characteristics, dealers may mark up bonds differently depending on the bond's characteristics. To the extent that this is

[28] In the Internet Appendix (Table IA.XII), we also examine OLS regressions to quantify the effects in a more traditional manner. Coefficient estimates may be difficult to interpret given all of the interaction effects identified by the machine learning model. Similar to the machine learning estimates, dealer fixed effects significantly increase the model's R^2, and dealer variables explain much of the R^2 improvement provided by dealer fixed effects.

true, it is desirable to allow bond and trade characteristics to interact with dealer variables. To do so, we predict markups based on dealer characteristics, rounding and strategic pricing indicator variables, bond-level trading activity, and the control variables from our previous regressions using the LightGBM gradient boosting decision tree library from Ke et al. (2017).

We use gradient boosting decision trees because, along with neural networks, they tend to be the best-performing machine learning method for tabular data, and decision trees have the additional advantage that they require little parameter tuning and are relatively easy to interpret (Chollet (2021), Howard and Gugger (2020)). The gradient boosting methodology builds decision trees based on successive binary splits of the independent variables to form "branches" of grouped training observations. For example, a split on bond maturity at 10 years would create a branch with two leaves: observations with maturities below 10 years and observations with maturities above 10 years. If those were the only two leaves, the predicted markup for bonds with maturities below 10 years would be the average of the markups of bonds with maturities of less than 10 years in the training sample. However, there will be many branches and leaves. The model grows leaf-wise, so after the first split it evaluates all leaves and variables and chooses the next leaf-variable combination that will improve the model the most. Nonlinearities are incorporated by splitting a single variable in multiple places, and interactions are incorporated by a succession of branches that split on different variables. The boosting component is the process whereby that once one tree is trained, the model residual errors are taken and a new tree is fit using those residuals as the dependent variable.

There are several potential model parameters to tune, including those related to the choice of split points, how many trees to train, and how many leaves to use per tree, among others. One of the benefits of tree-based methods is that they tend to be robust to these choices as we verify.[29]

C. Markup Determinants

Our main objective is to use dealer characteristics and practices to examine why some dealers charge higher markups than others. To evaluate how well dealer characteristics explain differences between dealers, Table V compares three models: *Base*, *Dealer Features*, and *Dealer Lag Variable*. The *Base* model includes bond and trade characteristics but not dealer characteristics. We estimate the model in a training period that includes all but the last year of data and then evaluate model performance out of sample in the last year of data. The *Dealer Features* model includes the *Base* model plus dealer characteristics and practices, and the *Dealer Lag Variable* model includes the *Base*

[29] We use the default parameters in our main analysis. In Table IA.XIII, we consider alternative hyperparameters and do not find much difference between the performance of models with the default parameters and the performance of models with optimized parameters found using Wang et al.'s (2021) tuning algorithm.

Table V

Performance of Gradient Boosting Decision Tree Models

This table reports R^2s for gradient boosting decision tree models predicting markups. The models are estimated using LightGBM and its default hyperparameter settings (Ke et al. (2017)). Models are estimated separately for new and seasoned issues in the training period and are then evaluated out of sample. The in-sample training period is July 2011 to December 2016 and the out-of-sample period is January 2017 to December 2017. The Base model uses bond and trade variables, which are the same as the green "Other Variables" in Figure 8. The Dealer Features model uses the base model variables plus the "Dealer-Level Variables" in Figure 8. The Dealer Lag model uses the base model variables plus the dealer's average markup on small customer purchases over the past 30 days.

	R^2	
	Training Period	Out of Sample
Panel A: New Issues		
Base	0.439	0.342
Dealer Features	0.527	0.427
Dealer Lag Var	0.521	0.454
Panel B: Seasoned Issues		
Base	0.543	0.479
Dealer Features	0.640	0.599
Dealer Lag Var	0.643	0.636

model plus the dealer's average markup over the prior 30 days. If the dealer variables well capture cross-sectional variation in dealer markups, then the *Dealer Features* model should perform almost as well in predicting markups as the dealer's past markups. We find that the *Dealer Features* model does indeed perform very well. The out-of-sample R^2 using the *Dealer Features* model is 43% (60%) for new (seasoned) issues compared to 45% (64%) using the *Dealer Lag Variable* model (Table V). The *Dealer Features* R^2 improvement over the *Base* model represents 76% of the improvement that the *Dealer Lag Variable* model achieves in both markets. We conclude that the dealer variables are able to explain most of the differences between dealers' average markups.

We next explore which variables in the *Dealer Features* model are most important for explaining markups and interpret their effects. To do so, we use Lundberg et al.'s (2020) methodology for identifying feature importance by calculating a variable's Shapley value. The output of this approach is known as a Shapley additive explanation (SHAP) value. For each trade, the method calculates the contribution of each variable to the model's predicted markup for the trade, including interaction effects. The sum of the SHAP values for all variables plus the average markup in the training sample equals the model's predicted markup for that trade. Following Lundberg et al. (2020), we calculate global variable importance as the average absolute value of a variable's SHAP values for all trades in the training sample.

Figure 8 plots average absolute SHAP values from the *Dealer Features*

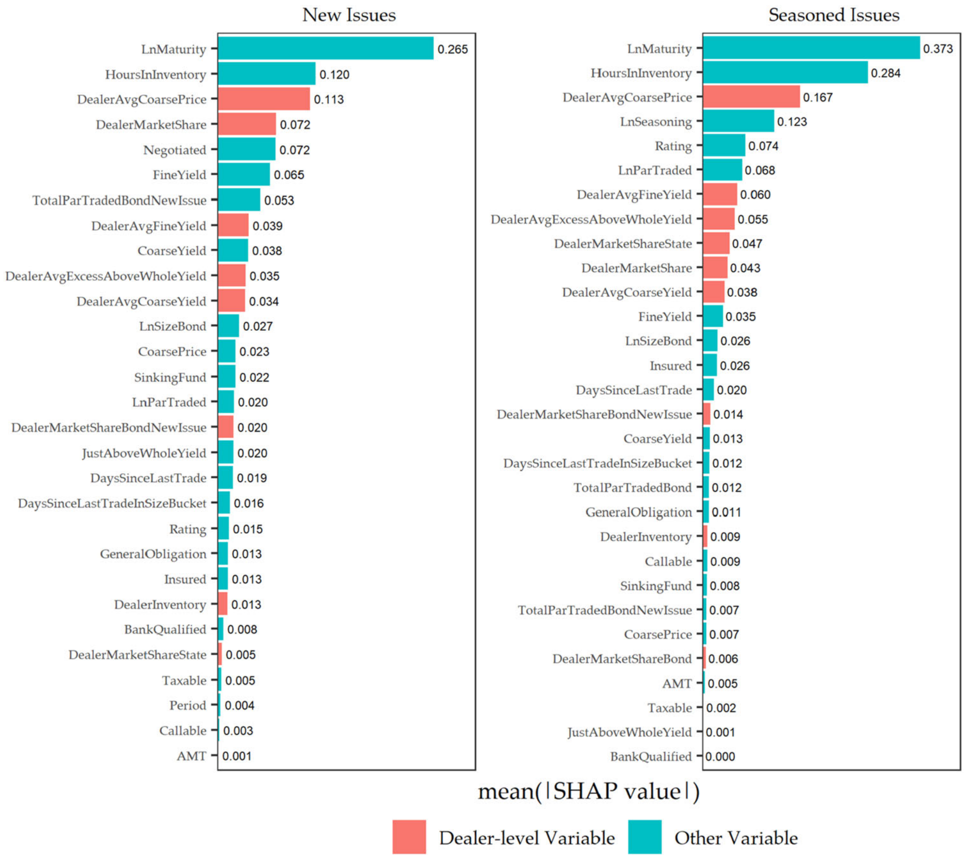

Figure 8. Average |SHAP| value for models predicting markups. Each plot shows the average absolute SHAP value for a variable calculated following Lundberg et al. (2020). The SHAP value for a variable is the estimated contribution of that variable to a trade's predicted markup. Predicted markups are estimated using the variables below and default parameters of the LightGBM gradient boosting decision tree model (Ke et al. (2017)). We estimate separate models for new and seasoned issues. The dependent variable is a trade's markup, and the sample is restricted to small trades from July 2011 to December 2016 to preserve the last year of data for out-of-sample testing. (Color figure can be viewed at wileyonlinelibrary.com)

model. The variables are sorted from most to least important. In both the new and seasoned issue markets, the top three explanatory variables are *LnMaturity*, *HoursInInventory*, and *DealerAvgCoarsePrice*. On average, *LnMaturity* explains 0.265 (0.373) ppt of the new (seasoned) issue predicted markup. The variable *HoursInInventory* measures how long the bond was in the dealer's inventory prior to the customer purchase transaction. Purchases that occur almost immediately after the dealer acquires the bond frequently have lower markups (Table IA.XIX), and this explains 0.120 (0.284) ppt of the new (seasoned) issue predicted markup. The variable *DealerAvgCoarsePrice* explains 0.113 (0.167) ppt of new (seasoned) issues. The model includes a trade-level coarse price dummy, so the effect of *DealerAvgCoarsePrice* captures how often dealers sell at coarse prices in general, not whether the specific trade is at a coarse price. The exact ordering of the remaining variables varies, but the dealer-level rounding and market share variables account for five of the top 11 new issue variables and six of the top 11 seasoned issue variables. In total, dealer-level variables contribute 0.331 (0.439) ppt to the predicted new (seasoned) issue markup (Figure IA.27). To shed light on factors associated with the highest markups, we also estimate a model with an indicator for markups above 3% as the dependent variable and find that the same factors associated with high average markups also predict extreme markups (Figure IA.29).

The average absolute SHAP values plotted in Figure 8 are useful for determining which variables have the most impact. To better understand when the variable's effect is positive or negative, we plot SHAP values across the range of individual variables with a focus on the dealer variables to help identify which dealer characteristics and pricing practices relate to differential markups. Figure 9 plots trade-level SHAP values for the two most important dealer pricing variables and the highest-impact dealer market share variable in each market.[30] Each dot in the plot is a trade. The x-axis shows the value of that variable for the trade and the y-axis shows the SHAP value—the estimated impact of the variable on that trade's markup. Trades with the same x-axis value can have different SHAP values due to interaction effects with other variables. The bar of dots at the far left of the plot shows SHAP values for trades for which the x-axis variable is not available. The gray histogram along the x-axis shows the distribution of the variable.

The results indicate that dealers who frequently execute trades at coarse prices tend to charge higher markups. The left plot in Panel A reports new issue results for *DealerAvgCoarsePrice* and the left plot in Panel B reports seasoned issue results. In both markets, the SHAP values increase from left to right across the x-axis, which corresponds to moving from dealers who rarely execute trades at coarse prices to dealers who frequently execute trades at coarse prices. SHAP values are particularly low for dealers who execute trades at coarse prices less than 10% of the time; the cluster of SHAP values of -0.5 or less indicates predicted markups from these dealers being 0.5 or more percentage points lower than average. A tendency to execute transactions at fine yields has the opposite relation to markups. The plots in the center column show that dealers who frequently trade at fine yields charge lower markups in both new and seasoned issues.

The two plots in the right column of Figure 9 show the effects of the most impactful dealer market share variables, *DealerMarketShare* in new issues and *DealerMarketShareState* in seasoned issues. In new issues, dealers with greater national market share charge higher markups. Similarly, in seasoned

[30] For completeness, Figure IA.28 includes SHAP value plots for the six variables with the largest average absolute SHAP values in new issues and in seasoned issues.

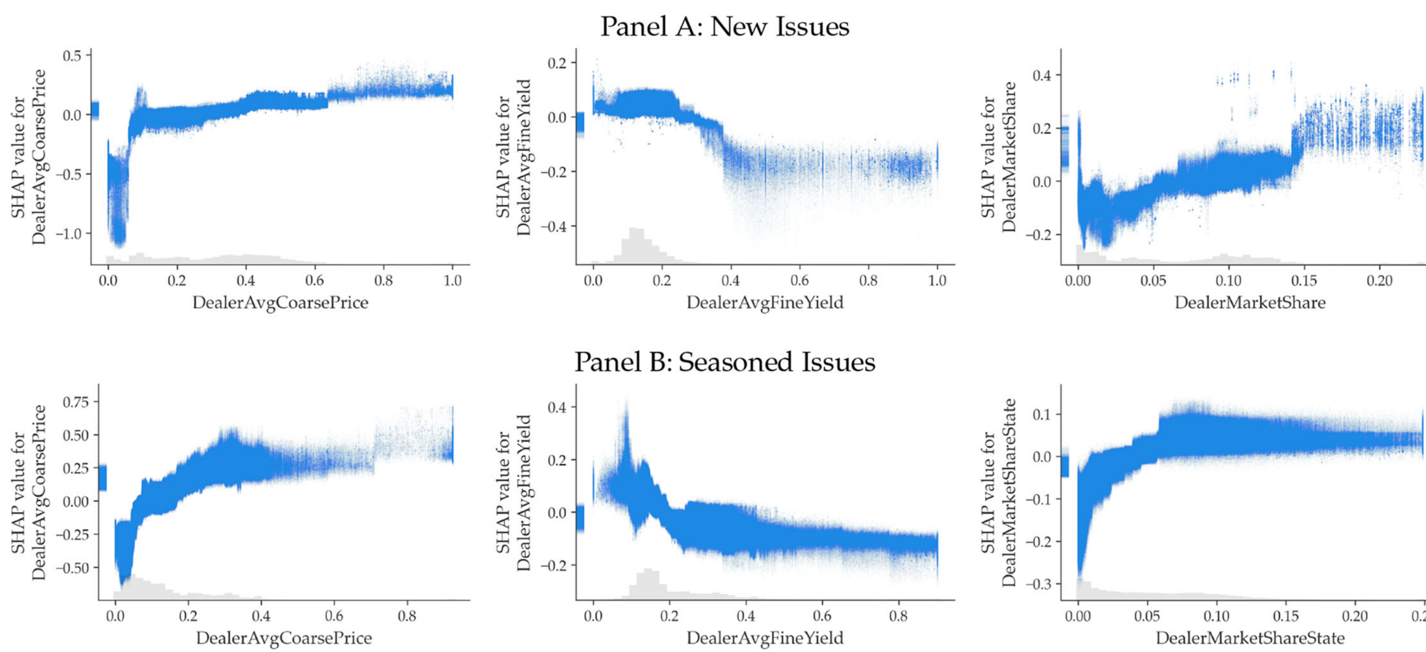

Figure 9. Effects of dealer characteristics on markups. This figure shows how individual dealer-level variables affect markups in the gradient boosting decision tree model. The models are the same as in Figure 8. They are estimated using the default parameters of the LightGBM gradient boosting decision tree model (Ke et al. (2017)). We report results for the main dealer-level rounding variables (*DealerAvgCoarsePrice* and *DealerAvgFineYield*) and for the dealer market share variable with the highest average absolute SHAP value in the new (seasoned) issue period. Each plot shows SHAP values across the range of a particular variable calculated following Lundberg et al. (2020). The SHAP value for a variable is the estimated contribution of that variable to a trade's predicted markup. Each dot is a trade. The x-axis shows the value of that variable for the trade and the y-axis shows the SHAP value—the estimated impact of the variable on the markup for that trade. Trades with the same x-axis coordinate can have different SHAP values (the y-axis) due to interaction effects with other variables. The vertical bar of dots at the far left of the plot shows SHAP values for trades for which the x-axis variable is not available. The gray histogram along the x-axis shows the distribution of the x-axis variable. (Color figure can be viewed at wileyonlinelibrary.com)

issues, trades from dealers with high market share in the state generally have higher markups.

Because it is the strongest predictor of markups in both the new and seasoned issue markets, we also examine SHAP values for maturity. For this analysis, we use the *Dealer Lag Variable* model to assess interactions between maturity and lagged dealer markups. Figure 10 plots the results. Each dot is a transaction, and the color of the dot represents the dealer's average lagged markup over the previous 30 days. Consistent with previous results, bonds with longer maturity tend to have higher markups, with this effect strongest for dealers with high lagged markups.

Overall, we find that our dealer-level rounding and market share variables are useful for predicting markups out of sample and are almost as good as knowing the dealer's actual average markup over the past 30 days. Even controlling for the pricing practices of a particular trade, one can predict whether a trade will have a high markup by knowing whether the dealer typically engages in pricing practices that exploit behavioral biases. Specifically, markups are high for dealers who frequently execute trades at coarse prices, rarely execute trades at fine yields, and have large market shares.

D. Strategic Pricing Determinants

We next assess what explains dealer pricing practices by estimating a model in which the dependent variable is an indicator for trades that have a coarse price, a coarse yield, or a yield that is within 0.5 bps below or 5 bps above a whole percentage-point yield. The model uses the same gradient boosting methodology as in the previous analysis. Dealer market share is the second-most important variable, and dealer average markup is the third-most important variable in both the seasoned issue and new issue markets.[31]

Figure 11 plots SHAP values for dealer average markup, dealer overall (national) market share, and the more granular market share that is most important in each model. The plots on the left show that dealers with high markups during the previous 30 days are more likely to use rounded or strategic pricing. In the middle column, rounding and strategic pricing decrease with dealer market share. This contrasts with the increasing markups observed as market share increases in Figure 9 and indicates that, all else equal, large dealers are less likely to engage in strategic pricing. However, relations are the opposite for more granular market share: in the new issue market, rounding and strategic pricing increase with bond-level market share, while in the seasoned issue market, rounding and strategic pricing increase with state-level market share.

[31] Trading volume is the most important variable in the new issue market, and hours in inventory is the most important variable in the seasoned issue market. See Figure IA.30 for average absolute SHAP values for all variables included in the models.

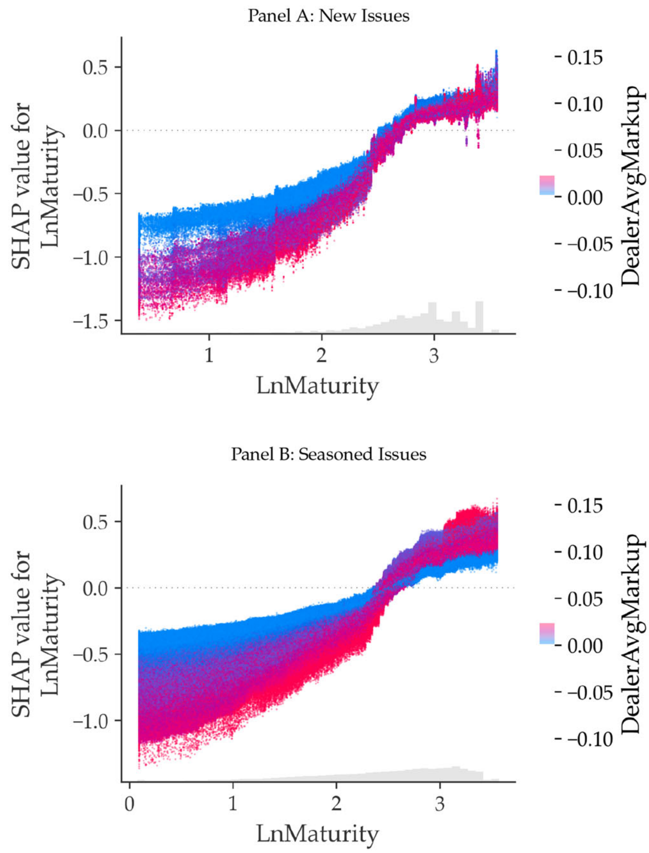

Figure 10. Effect of maturity on markups, interacted with a dealer's lag average markup. This figure shows how maturity affects markups in the gradient boosting decision tree models that use the base model variables (the green variables in Figure 8) plus a dealer's lag average markup over the prior 30 days to predict markups. The models are estimated using the training sample, July 2011 to December 2016. It also shows how the maturity effect interacts with a dealer's average markup over the past 30 days. Models for new and seasoned issue markets are estimated separately using the default parameters of the LightGBM gradient boosting decision tree model (Ke et al. (2017)). Each plot shows SHAP values across the range of a particular variable calculated following Lundberg et al. (2020). The SHAP value for a variable is the estimated contribution of that variable to a trade's predicted markup. Each dot is a trade. The x-axis shows the value of that variable for the trade and the y-axis shows the SHAP value—the estimated impact of the variable on the markup for that trade. Trades with the same x-axis coordinate can have different SHAP values (the y-axis) due to interaction effects with other variables. The vertical bar of dots at the far left of the plot shows SHAP values for trades for which the x-axis variable is not available. The gray histogram along the x-axis shows the distribution of the x-axis variable. The color of each dot shows the dealer's lagged average markup to visualize interactions. Red indicates high values for the dealer's lagged average markup and blue indicates low values. (Color figure can be viewed at wileyonlinelibrary.com)

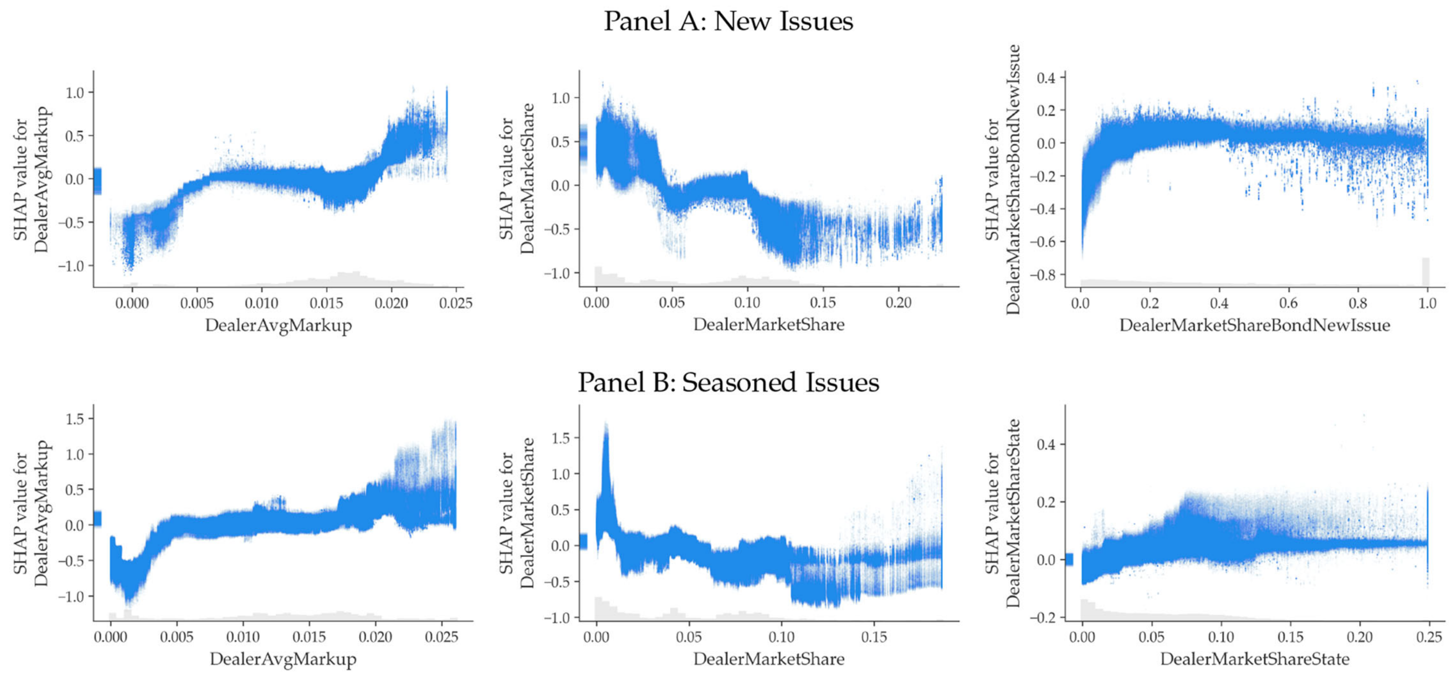

Figure 11. Effects of dealer characteristics on strategic pricing. This figure shows how individual dealer-level variables affect strategic pricing in the gradient boosting decision tree model. We report results for the dealer's average markup during the past 30 days and for the two dealer market share variables with the highest average absolute SHAP value in the new (seasoned) issue period. The dependent variable is a dummy variable that equals one if the trade is executed at a coarse price, a coarse yield, or a yield just above a whole yield and zero otherwise. The independent variables are the base model variables (the green variables in Figure 8) plus dealer market share variables and the dealer's average markup over the prior 30 days. The full set of variables is reported in Figure IA.30. The models are estimated using the training sample, from July 2011 to December 2016. The models are estimated separately for new and seasoned issues using the default parameters of the LightGBM gradient boosting decision tree model (Ke et al. (2017)). Each plot shows SHAP values across the range of a particular variable calculated following Lundberg et al. (2020). The SHAP value for a variable is the estimated contribution of that variable to a trade's predicted probability of strategic pricing. Each dot is a trade. The x-axis shows the value of that variable for the trade and the y-axis shows the SHAP value—the estimated impact of the variable on strategic pricing for that trade. Trades with the same x-axis coordinate can have different SHAP values (the y-axis) due to interaction effects with other variables. The vertical bar of dots at the far left of the plot shows SHAP values for trades for which the x-axis variable is not available. The gray histogram along the x-axis shows the distribution of the x-axis variable. (Color figure can be viewed at wileyonlinelibrary.com)

VII. MSRB Rule Changes, Markup Quantification, and Other Markup Explanations

In this section, we examine the impact of MSRB rule changes on markups, the total dollar impact of small trade markup practices, and other potential explanations for markup variation.

A. Are MSRB Regulatory Changes Effective?

We begin by assessing changes around two MSRB rules changes and then discuss details and implementation. Beginning on March 21, 2016, the MSRB implemented a best execution rule requiring dealers to attempt to find the best possible price for municipal bond transactions. On May 14, 2018, the MSRB started to require that dealers disclose additional cost and markup information to their customers. Did this additional regulation and transparency affect markups? The time-series evidence in Figure 1 shows little effect of either rule change in the new or seasoned issue market, and markups are also unchanged in narrower windows around the rule changes (Figure IA.33).

Event regressions of trade-level markups within six weeks before and after each event on trade and bond characteristics plus an indicator for the trade occurring after the rule change are reported in Table VI.[32] For the best-execution rule change, the *Post* coefficient is not significantly different from zero in the new issue market (column (1)) and slightly negative (−2.6 bps) in the seasoned issue market (column (3)), and both point estimates are small relative to the average small-trade markups of 1.33% in the new issue market and 0.946% in the seasoned issue market during the event period. We also test for a reduction in markups over 3% and find no significant effect in the new issue market and a slight increase (with 10% significance) in the seasoned issue market.[33] For the disclosure event, we find no impact in the new issue market and a slight reduction in seasoned issue markups of −5.6 bps relative to mean markups of 1.07 ppt (column (7)).[34] The incidence of extreme markups over 3% falls more meaningfully by a statistically significant 1.2 ppt (from 5.7% to 4.5%) in the seasoned issue market but rises by an insignificant 1.4 ppt in the new issue market (columns (8) and (6)).

Why have the transparency, fair pricing, and best-execution regulations not sizably reduced markups? One limitation of the MSRB transparency and regulatory protections changes is that municipal bond trading is a retail over-the-counter market. Electronic trading has been introduced between

[32] Covariates are similar before and after the rule changes, and regression results are similar in univariate specifications with no control variables (Table IA.XV).

[33] We also explore whether the limited effect could be due to high-cost dealers being less likely to comply with regulation (Table IA.XIV), but we do not find a statistically significant difference between high- and low-cost dealers' markup changes surrounding the rule changes. This analysis is not possible for the disclosure event because it occurred after the end of our anonymized dealer ID data.

[34] Cuny, Even-Tov, and Watts (2021) find that markups decreased similarly (−4.4 bps) following the same rule in the corporate bond market.

Table VI

Disclosure and Best Execution Events

This table reports regressions of small (≤\$100k) trade markups on *Post* event dummies indicating trades occurring after two regulatory changes. We include trades from six weeks before to six weeks after each event. The dependent variable is either the markup or a dummy variable for markups over 3%. Regressions for both dependent variables are estimated using OLS. The best execution event (March 21, 2016) is the day the MSRB best execution rule went into effect. The markup disclosure event (May 14, 2018) is the day brokers were first required to disclose markups to customers. We include customer purchases only. New issue trades occur from the start of trading to 14 days later, and seasoned issue trades occur more than 90 days after the bond's issue date. The best execution event falls within our dealer-level sample, so we use dealer-level data to calculate markups. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. Matches between a customer trade and the dealer's purchase price are made on a FIFO basis. We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. The disclosure event is after our dealer-level data end, so we use the WRDS data set without dealer IDs. The new issue markup is the percentage premium over the offering price. The seasoned issue markup is the estimated total spread captured by the dealer sector when matching customer sales to eventual customer purchases. The markup is (Customer Purchase Price/Matched Dealer Cost) − 1. The seasoned issue matches consist of the combined set of immediate, round-trip, and FIFO matches calculated following Green, Hollifield, and Schürhoff (2007b). We limit the seasoned issue sample to matches with no more than seven days between matched purchases and sales. Both new and seasoned issue regression control variables include ln(trade's par amount), ln(bond's maturity), ln(bond's total par issued), and dummy variables indicating the bond's credit rating on a scale from 0 (unrated) to 24 (AAA), the bond's state, insured bonds, general obligation bonds, callable bonds, bonds with a sinking fund, bank qualified bonds, taxable bonds, and AMT bonds. In addition, new issue bonds have a dummy for negotiated offerings and dummies for each day since the bond was offered to the public, while seasoned issues add a control for ln(seasoning in days since issuance). All variables are winsorized at the 0.5% and 99.5% levels. Three-way clustered (dealer, bond, and day) standard errors are in parentheses. Significance levels are indicated by asterisks for *p*-values less than or equal to 0.01 (***), 0.05 (**), and 0.1 (*).

	Best Execution Event				Disclosure Event			
	New Issues		Seasoned Issues		New Issues		Seasoned Issues	
$y = (M$ is Markup)	M (1)	$M > 3\%$ (2)	M (3)	$M > 3\%$ (4)	M (5)	$M > 3\%$ (6)	M (7)	$M > 3\%$ (8)
Post	0.008	0.001	-0.026^{***}	0.004^{*}	−0.001	0.014	-0.056^{***}	-0.012^{***}
	(0.026)	(0.003)	(0.006)	(0.002)	(0.062)	(0.014)	(0.011)	(0.002)
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Days Since Offering FE	✓	✓			✓	✓		
State and Rating FE	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-level Data?	✓	✓	✓	✓				
Observations	126,651	126,651	255,514	255,514	73,330	73,330	255,403	255,403
Adjusted R^2	0.334	0.023	0.278	0.044	0.417	0.128	0.301	0.076
Dependent Variable Mean	1.33	0.024	0.946	0.018	1.10	0.056	1.07	0.051

dealers and is gaining traction (Cestau et al. (2018)), but pretrade price information based on available quotes is not available to retail investors (Craig, Kim, and Woo (2018), Wu, Bagley, and Vieira (2018)). This makes it difficult for a customer to ascertain whether they are receiving best execution.[35] Limited

[35] Industry advocates pointed this out in a comment letter to the MSRB: "there are very rarely numerous buyers and sellers of a given municipal security over a short time span, which would

market information also complicates markup disclosure because it is difficult to determine the fair market price of a security. Perhaps as a result, the current MSRB markup disclosure rule (MSRB Rule G15(a)(i)(F)) applies only to retail transactions in which the dealer has an offsetting trade on the same day. An MSRB report (Wu (2018)) estimates that this applies to less than 40% of retail trades.

A second limitation is that while traders have access to nearly real-time transaction price information through MSRB's EMMA website, it is not clear whether most retail customers know how to access and interpret this information.[36] If customers effectively use EMMA transaction data, they will have better information on prices when there are more recent trades. To examine this possibility, we analyze how markups vary with recent trading activity. We find no evidence that markups are lower in periods of high trading activity (Figure IA.34).[37]

A third limitation of the MSRB regulation is that enforcement appears to be very limited. Between August 2016 and August 2021, the Financial Industry Regulatory Authority (FINRA) took only 12 disciplinary actions involving 204 transactions related to fair pricing of municipal bonds (Table IA.XVI), and these actions typically correspond only to the most egregious examples of unfair pricing. Discussions within the FINRA decisions do not provide a clear definition of what constitutes an unfair markup, but enforcement generally appears to be limited to markups above 3%.[38] In addition, 204 disciplinary transactions over the course of five years is small relative to 90,000 purchases of any trade size per year in the transaction data with markups over 3%. Even with this high threshold, enforcement appears spotty. This finding lines up

(if they existed) allow market price convergence to occur and permit an equity-type best execution rule to be meaningful" (SIFMA comment letter to the MSRB, March 13, 2014). Related, NYSE Euronext noted that "the MSRB should propose a rule that will advance the efforts of pre-trade price transparency" (NYSE Euronext comment letter the MSRB, March 31, 2014).

[36] For example, the EMMA website requires a Committee on Uniform Security Identification Procedures (CUSIP) code without a link to look up CUSIPs, something seldom required for stocks. The Consumer Federation of America argues that "It's not realistic to expect retail investors to use TRACE and EMMA with any reasonable degree of expertise. In order to use TRACE and EMMA, one has to know each website exists and what specifically each website offers. It would likely confuse an investor that he or she has to go to different websites to see different types of recent bond transactions… Finally, assuming that a retail investor knows what information to look for and finds it, one would need to be able to understand and make use of that information for one's benefit" (Consumer Federation of America comment letter to the MSRB, January 20, 2015).

[37] Before controlling for dealer fixed effects, markups are higher for bonds that traded recently (see the top row of Figure IA.34), which is the opposite of what availability of recent transaction information in EMMA would predict. Recent trading activity is also unimportant in the machine learning models predicting markups (Figures 8 and IA.29). Recent trading activity is somewhat predictive of strategic pricing in the seasoned issue market (Figure IA.30), but bonds with recent trading activity have moderately higher usage of strategic pricing (Figure IA.32), again inconsistent with what one would expect from more information in EMMA.

[38] For example, "We were unable to locate an earlier case in which a court, FINRA, or the SEC held that markups below 3 percent were in violation of MSRB Rule G-30" (FINRA Disciplinary Proceeding 20120317482-03, September 26, 2017).

with Macey and O'Hara's (1997) warning that best-execution rules are unlikely to be effective.

B. What is the Aggregate Magnitude of Excess Markups?

Total markup revenue over the 6.5-year sample with dealer-level data is \$1.1 billion in the new issue market and \$2.9 billion in seasoned issue market (Table IA.XVII). Because there is no consensus in the literature on how to define excessive markups, we provide a range of estimates rather than take a hard stance on what markups ought to be. If markups over 0.5% were reduced to 0.5% (which is approximately the average markup for medium and large trades), customers would have saved \$840 million on new issue markups and \$2.0 billion on seasoned issue markups. With a higher cap of 1% (2%), customers would have saved \$550 (\$130) million in new issues and \$1.3 billion (\$349 million) in seasoned issues. Despite declines in seasoned issue markups, revenue from the portion of markups over 1% still accounted for 45% (36%) of dealers' new (seasoned) issue markup revenue in 2017. Excess markups associated with rounded prices and nonfine yields during the 6.5-year sample period are estimated to be \$246 million in the new issue market and \$624 million in the seasoned issue market, which is 21% of total markup revenue in both markets.

Overall, extreme markup costs are large relative to total markups. Since our sample and analysis is limited to customer purchase transactions, we do not capture profits when a dealer buys a bond from a customer and sells it to another dealer. The total effect would be larger including these markdowns.

C. Are There Common Linkages between the New Issue and Seasoned Issue Markets across Dealers?

Thus far we have thus far examined a dealer's pricing of new and seasoned issues separately. We now test additional explanations for dealer markups by jointly analyzing the two markets.

Search costs are much more important in the seasoned issue market than in the new issue market, where shares are allocated directly by the underwriter. If cross-sectional differences in new and seasoned issue markups represent cost differences unique to one of the markets or other random variation unrelated to specific dealers, markups need not be correlated in the two markets. However, if differences across dealers represent different approaches to customer pricing, a dealer's markups should be correlated in the two markets. This is what we find in Figure IA.35.[39]

Another potential connection between the new and seasoned issue markets is that customers may be willing to pay more for a bond in the new issue mar-

[39] Several brokers consistently have both new issue and seasoned issue markups of 10 bps or less, and there is a larger mass of brokers that have markups above 1.5% in both the new issue and seasoned issue markets. The correlation between median dealer markups in the new issue and seasoned markets is 0.535 and highly significant.

ket if a dealer will buy it back at a favorable price in the seasoned issue market. In contrast, new issue purchase markups and seasoned issue sale markdowns have a positive and significant correlation of 0.438 (Figure IA.36).

Dealers who play a large role in the new issue market, perhaps through their role in underwriting, frequently continue to facilitate trade in the bond in the seasoned issue market (Figure IA.37 and Table IA.XVIII). To assess the relation between new issue market share and seasoned issue pricing, we reestimate our gradient boosted model using only seasoned issue trades for which new issue market share is available.[40] In this restricted sample, new issue market share is the fifth-most important variable for explaining seasoned issue markups (Figure IA.38, Panel A). In addition, dealers who had high market share when a bond was issued, and thus were likely part of the bond's syndicate, typically charge higher seasoned issue markups (Figure IA.38, Panel B).

D. Inventory and Competition

Markups vary with the amount of time a bond is held in a dealer's inventory.[41] However, time in inventory does not explain markup variance across dealers, and dealers with high markups on inventory held less than a day also have high markups on inventory held for longer periods (Figure IA.39).

We also consider how competition affects markups by regressing markups on the number of dealers trading the bond (Table IA.XX). Instead of dampening markups, competition is associated with moderately higher markups. This suggests that retail investors obtain little benefit from being able to access a bond from multiple dealers, which is consistent with what we would expect if most retail customers have relationships only with a single dealer or a small number of dealers. Lack of customer sophistication may be one reason for high and varied markups but this is difficult to capture empirically.[42]

[40] We included a dealer's new issue market share for a bond in the models predicting seasoned issue markups in Figure 8. However, that figure obscures *DealerMarketShareBondNewIssue*'s importance because the variable is missing for all bonds issued before our sample starts.

[41] Empirically, markups are lowest when dealers purchase and sell a bond within one minute, which represents 32% of small seasoned issue trades (Table IA.XIX). Bonds sold out of inventory that has been held for at least a day have average markups that are over 1 ppt higher. This potentially compensates dealers for inventory risk, but the actual risk is likely very low because even bonds that are held in inventory are typically sold within one to five days. Recall that our sample was restricted to inventory held seven days or less.

[42] In Tables IA.XXI and IA.XXII, we consider two potential proxies: investors buying a bond because other bonds by the issuer have recently matured may be less knowledgeable than investors who trade more frequently (Table IA.XXI), and investors in high-tax states where there is a greater incentive to invest in municipal bonds may also be less knowledgeable (Table IA.XXII). In both cases, there is some evidence of higher markups, but the effects are inconsistent and economically small. Dealer network centrality could also affect markups, but we find that network centrality has little explanatory power for markup variation across dealers (see Section I of the Internet Appendix).

VIII. Conclusion

Unlike equities, municipal bonds trade over the counter in a dealer market that is opaque to the many retail investors who comprise a large percentage of the market. These investors often rely on their brokers to provide them with sound investing advice and fair pricing, and that is what the industry and MSRB rules purport to provide. We empirically investigate whether municipal bond dealers deliver on their commitment to best execution and fair pricing. We find skeptical results.

Instead of delivering uniform pricing, dealer transactions with customers occur at highly variable markups relative to both reoffering prices and dealer costs. On the same day, customers frequently buy the same bond at different prices from different dealers, and prices even vary across different customers purchasing the same bond from the same dealer on the same day. These price differences are not explained by trade characteristics or by dealer costs. Some dealers provide customers low and consistent markups, but this does not appear to be the industry norm. Pricing at quarter or eighth price or yield increments is common and is seemingly a method to obtain higher markups. High-cost dealers also appear to use strategic pricing around yield thresholds to disguise higher markups, and markups are higher for long-maturity bonds where their impact on yields is less visible. Machine learning shows that a dealer's past strategic pricing practices are strong predictors of a bond's markup. It would be useful to know the specific dealers who most frequently engage in these practices, but this is not possible due to data anonymity requirements. Additional disclosure along these lines may be useful for investors.

A common theme across all of these practices is that dealers appear to use their pricing discretion to charge higher markups to small customers when investors are less likely to notice. More broadly, our findings raise concerns that conflicts of interest may be widespread in the financial services industry. Since most major financial firms that trade municipal bonds also sell other products with more limited transparency, our findings suggest that more transparency, investor protection, and academic research is needed to further examine the "rent-seeking" dimension of finance (Zingales (2015)).

Initial submission: May 6, 2020; Accepted: May 27, 2022
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Babina, Tania, Chotibhak Jotikasthira, Christian T. Lundblad, and Tarun Ramadorai, 2021, Heterogeneous taxes and limited risk sharing: Evidence from municipal bonds, *Review of Financial Studies* 34, 509–568.

Bergstresser, Daniel, and Randolph Cohen, 2015, Changing patterns in household ownership of municipal debt: Evidence from the 1989-2013 surveys of consumer finance, Working paper, Brandeis University.

Bergstresser, Daniel, and Patrick Herb, 2022, Do risk premia explain dealer markups in municipal bond offerings? Working paper, Brandeis University.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.
Bessembinder, Hendrik, and William Maxwell, 2008, Markets: Transparency and the corporate bond market, *Journal of Economic Perspectives* 22, 217–234.
Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities, and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251–288.
Bessembinder, Hendrik, Chester Spatt, and Kumar Venkataraman, 2020, A survey of the microstructure of fixed-income markets, *Journal of Financial and Quantitative Analysis* 55, 1–45.
Bhattacharya, Utpal, Craig W. Holden, and Stacey Jacobsen, 2012, Penny wise, dollar foolish: Buy–sell imbalances on and around round numbers, *Management Science* 58, 413–431.
Biais, Bruno, and Richard Green, 2019, The microstructure of the bond market in the 20th century, *Review of Economic Dynamics* 33, 250–271.
Carslaw, Charles A.P. N., 1988, Anomalies in income numbers: Evidence of goal-oriented behavior, *Accounting Review* 3, 321–327.
Cestau, Dario, Burton Hollifield, Dan Li, and Norman Schürhoff, 2018, Municipal bond markets, *Annual Review of Financial Economics* 11, 65–84.
Chalmers, John, Steve Liu, and Zhi Jay Wang, 2021, The difference a day makes: Timely disclosure and trading efficiency in the muni market, *Journal of Financial Economics* 139, 313–335.
Chollet, Francois, 2021, Deep Learning with Python (Simon and Schuster, New York, NY).
Christie, William G., and Paul H. Schultz, 1994, Why do NASDAQ market makers avoid odd-eighth quotes?, *Journal of Finance* 49, 1813–1840.
Craig, Louis, Abby Kim, and Seung Won Woo, 2018, Pre-trade information in the municipal bond market, Technical report, Securities and Exchange Commission.
Cuny, Christine, Omri Even-Tov, and Edward M. Watts, 2021, From implicit to explicit: The impact of disclosure requirements on hidden transaction costs, *Journal of Accounting Research* 59, 215–242.
Dimmock, Stephen G., William C. Gerken, and Nathaniel P. Graham, 2018, Is fraud contagious? Coworker influence on misconduct by financial advisors, *Journal of Finance* 73, 1417–1450.
Duflo, Esther, 2017, The economist as plumber, *American Economic Review* 107, 1–26.
Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–1451.
Egan, Mark, 2019, Brokers vs. retail investors: Conflicting interests and dominated products, *Journal of Finance* 74, 1217–1260.
Egan, Mark, Gregor Matvos, and Amit Seru, 2019, The market for financial adviser misconduct, *Journal of Political Economy* 127, 233–295.
Ellison, Glenn, 2006, Bounded Rationality in Industrial Organization, Vol. 2 of *Econometric Society Monographs*, 142–174 (Cambridge University Press, Cambridge, UK).
Gabaix, Xavier, and David Laibson, 2006, Shrouded attributes, consumer myopia, and information suppression in competitive markets, *Quarterly Journal of Economics* 121, 505–540.
Garrett, Daniel, 2021, Conflicts of interest in municipal bond advising and underwriting, Working paper, University of Pennsylvania.
Goldstein, Michael A., Edith S. Hotchkiss, and Stanislava Nikolova, 2021, Dealer behavior and the trading of newly issued corporate bonds, Working paper, Boston College.
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.
Goldstein, Michael A., Shane M. Moser, and Robert A. Van Ness, 2022, Trade price clustering in the corporate bond market, *China Finance Review International* 12, 353–377.
Green, Richard C., 2007, Presidential address: Issuers, underwriter syndicates, and aftermarket transparency, *Journal of Finance* 62, 1529–1550.
Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007a, Dealer intermediation and price behavior in the aftermarket for new bond issues, *Journal of Financial Economics* 86, 643–682.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, 2007b, Financial intermediation and the costs of trading in an opaque market, *Review of Financial Studies* 20, 275–314.

Green, Richard C., Dan Li, and Norman Schürhoff, 2010, Price discovery in illiquid markets: Do financial asset prices rise faster than they fall?, *Journal of Finance* 65, 1669–1702.

Harris, Larry, 2015, Transaction costs, trade throughs, and riskless principal trading in corporate bond markets, Working paper, University of Southern California.

Harris, Larry, Albert S. Kyle, and Erik R. Sirri, 2015, Statement of the Financial Economists Roundtable, April 2015: The structure of trading in bond markets, *Financial Analysts Journal* 71, 5–8.

Harris, Lawrence, 1991, Stock price clustering and discreteness, *Review of Financial Studies* 4, 389–415.

Harris, Lawrence E., and Michael S. Piwowar, 2006, Secondary trading costs in the municipal bond market, *Journal of Finance* 61, 1361–1397.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, 2021, Discriminatory pricing of over-the-counter derivatives, *Management Science* 67, 6660–6677.

Hendershott, Terrence, and Ananth Madhavan, 2015, Click or call? Auction versus search in the over-the-counter market, *Journal of Finance* 70, 419–447.

Howard, Jeremy, and Sylvain Gugger, 2020, Deep Learning for Coders with Fastai and PyTorch: AI Applications without a PhD (O'Reilly Media, Sebastopol, CA).

Kahn, Charles, George Pennacchi, and Ben Sopranzetti, 1999, Bank deposit rate clustering: Theory and empirical evidence, *Journal of Finance* 54, 2185–2214.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu, 2017, LightGBM: A highly efficient gradient boosting decision tree, *Advances in Neural Information Processing Systems* 30, 3146–3154.

Li, Dan, 2007, New Findings in the Market Microstructure of Over-the-Counter Markets, Ph.D. thesis, Carnegie Mellon University.

Li, Dan, and Norman Schürhoff, 2019, Dealer networks, *Journal of Finance* 74, 91–144.

Lundberg, Scott M., Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M. Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee, 2020, From local explanations to global understanding with explainable AI for trees, *Nature Machine Intelligence* 2, 2522–5839.

Macey, Jonathan R., and Maureen O'Hara, 1997, The law and economics of best execution, *Journal of Financial Intermediation* 6, 188–223.

Schultz, Paul, 2001, Corporate bond trading costs: A peek behind the curtain, *Journal of Finance* 56, 677–698.

Schultz, Paul, 2012, The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors, *Journal of Financial Economics* 106, 492–512.

Shiller, Robert J., 2000, Irrational Exuberance (Princeton University Press, Princeton, NJ).

Thomas, Manoj, and Vicki Morwitz, 2005, Penny wise and pound foolish: The left-digit effect in price cognition, *Journal of Consumer Research* 32, 54–64.

Wang, Chi, Qingyun Wu, Markus Weimer, and Erkang Zhu, 2021, FLAML: A fast and lightweight AutoML library, MLSys Proceedings.

Wu, Simon Z., 2018, Transaction costs for customer trades in the municipal bond market: What is driving the decline? Technical report, Municipal Securities Rulemaking Board.

Wu, Simon Z., John Bagley, and Marcelo Vieira, 2018, Analysis of municipal securities pre-trade data from alternative trading systems, Technical report, Municipal Securities Rulemaking Board.

Wu, Simon Z., and Marcelo Vieira, 2019, Mark-up disclosure and trading in the municipal bond market, Technical report, Municipal Securities Rulemaking Board.

Zingales, Luigi, 2015, Presidential address: Does finance benefit society?, *Journal of Finance* 70, 1327–1363.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication Code.
Disclosure Statement.