

The Erosion of Retail-Distribution Rents in Municipal Bond Underwriting

John Hund, Christian Lundblad, Christos A. Makridis, and Giang Nguyen*

September 22, 2026

Abstract

We study the decline in municipal bond underwriters' compensation amid rising transparency and investor sophistication. Using new-issue data from 2005–2024, we estimate a structural model of markups with latent classification of investor type. We find a sharp decline in both the probability of uninformed trades and the markups underwriters charge on them. As a result, underwriters' distribution rent falls from 15.99 to 3.71 basis points of issuance (a 76.8% decline), while underwriting spreads remain flat, indicating sustained pressure on underwriter compensation. Using the model estimates, we find that greater transparency is associated with a lower probability of uninformed trades, while higher household financial capability is associated with lower markups on those trades. Finally, we show that the conventional \$100,000 trade-size threshold performs poorly as a proxy for investor informedness.

Keywords: Public Finance; Municipal Bonds; Underwriters; Markup; Transparency

JEL Codes: G14, G23, G24

*John Hund: University of Georgia, email: jhund@uga.edu; Christian Lundblad: University of North Carolina at Chapel Hill, email: christian_lundblad@unc.edu; Christos Makridis: University of Nicosia and Stanford University, email: cmakridi@stanford.edu; Giang Nguyen: Pennsylvania State University, email: gxnl3@psu.edu. We thank conference participants and seminar audience members at the 13th Annual Brookings Municipal Finance Conference, the 2025 SEC Conference on Financial Market Regulation, the University of Texas (Austin) Finance Alumni Conference, Penn State University, University of Denver, the Federal Reserve Board of Governors, Imperial College London, and Virginia Tech University for valuable comments. This paper was previously titled “Investor Sophistication and Municipal Underwriting” and, before that, “Markup or Markdown: National Underwriters’ Exit and the Changing Landscape of Municipal Finance.” All errors and the views expressed in this paper are our own and do not reflect those of any affiliated institutions.

1 Introduction

The \$3.8 trillion municipal bond market is a vital source of capital for financing public infrastructure and services across the United States. Municipal underwriters play a central role in this ecosystem: they organize bond issuance, purchase securities from issuers, and resell them to investors, earning compensation through underwriting spreads collected from issuers and markups charged on customer trades. Over the past two decades, however, the structure of this market has changed markedly for at least two reasons. First, regulatory initiatives have increased price transparency (Wu, 2018). Second, the investor base has become increasingly sophisticated, with mutual funds, ETFs, and separately managed accounts displacing a historically retail-dominated clientele (Cestau et al., 2019). These developments have important implications for the economics of underwriting, especially for the rents underwriters can earn from trading with uninformed investors in the new issue market.

The conceptual framework of Green (2007) provides a natural lens for understanding these developments. In this model, dealer-underwriters behave as Bertrand competitors in the primary market, which keeps underwriting spreads tightly constrained. Underwriter compensation therefore stems largely from markups charged to uninformed investors in the new issue market—markups enabled by limited price transparency and heterogeneous information-gathering capabilities among market participants. As transparency rises and the share of retail investors declines, the model predicts a contraction in underwriters’ ability to extract such rents. Yet despite the centrality of this mechanism, there is limited empirical evidence on whether and to what extent underwriter compensation has actually evolved in response

to these long-run changes in market microstructure.

Our study aims to fill this gap. While prior work documents generally declining municipal bond transaction costs (e.g., [Schultz, 2012](#); [Chalmers et al., 2021](#); [Wu, 2018](#); [Griffin et al., 2023](#)), the evidence is mixed for new-issue markups, which are central to the underwriter business model. Moreover, the literature has relied on the conventional \$100,000 trade-size threshold to differentiate retail (presumably uninformed) from institutional (presumably informed) investors (e.g., [Hong and Warga, 2004](#); [Harris and Piwowar, 2006](#); [Green et al., 2006](#); [Schultz, 2012](#); [Cuny, 2017](#); [Hendershott et al., 2020](#), as well as SEC Rule 15c2-12 analyses in [Ivanov et al., 2024](#)). This proxy, however, has not been evaluated in light of significant changes in transparency and investor composition over the past two decades.

We answer these questions by assembling a comprehensive database integrating the Mergent Municipal Fixed Income Securities database, MSRB trade data, the SDC Platinum database, and Bloomberg for all municipal bonds issued since 2005 until the end of 2024 (data cutoff point). Over this 20-year sample period, we document that underwriting spreads and new-issue markups have declined substantially over time, especially among national underwriters who account for the majority of underwriting activity. These time-series trends persist after controlling for a wide range of bond, issuance, and issuer characteristics.

To understand the mechanism behind these declines, we estimate a structural model grounded in [Green \(2007\)](#) and [Green et al. \(2007\)](#). The model links observed markups in the new issue market to bond, issue, and trade characteristics, but permits the markup schedule to differ for informed and uninformed investors. Because investor type is not directly observable, we follow [Green et al. \(2007\)](#) in modeling informedness as a latent variable that depends on trade, bond, and issue characteristics (including underwriting spreads), reflecting

investors' endogenous decision to incur information-acquisition costs. The empirical framework is a mixture of regression models consisting of an outcome equation for markups and a two-regime classification equation for investor type.

Our estimates reveal two striking patterns. First, the probability that underwriters encounter uninformed investors has fallen sharply, from more than 57% in 2006 to 26% in 2024. Second, the average markup charged to uninformed investors has declined from roughly 100 basis points (bps) in 2006 to below 80 bps in 2024, after peaking at 152 bps during the 2009 financial crisis. Together, these trends imply that underwriters are now half as likely to transact with uninformed customers as in 2006 and charge them roughly half of what they did at the 2009 peak. We combine these margins into a single object, the retail-distribution rent per dollar of issuance, and find that it falls from 15.99 bps in 2006 to 3.71 bps in 2024, a 76.8% decline. A Shapley decomposition attributes 61.9% of the decline to lower uninformed order flow and 38.1% to smaller markups on that flow. Underwriting spreads do not rise to offset the loss, so the rent's share of total underwriter compensation falls from 22.27% to 8.15%.

Using the structural estimates of the prevalence of uninformed order flow and the markup wedge between uninformed and informed trades, we examine the effects of rising transparency and greater investor sophistication. We treat the May 2018 implementation of MSRB Rules G-15 and G-30, which mandated markup disclosure, as a transparency shock. After the shock, the probability of uninformed trading falls significantly, but the terms received by trades that remain uninformed do not improve. We then exploit variation in investor sophistication across states and over time, in both direct and delegated forms. Direct sophistication is associated with a narrower markup wedge between the uninformed and informed

regimes, suggesting that sophisticated households’ option to become informed disciplines dealer pricing even when it is not exercised. Overall, public disclosure reduces the prevalence of uninformed trading, whereas direct household sophistication lowers the price paid by those who remain uninformed.

Finally, our results allow us to re-evaluate the widely used \$100,000 trade-size cutoff for distinguishing between retail and institutional investors. We find that trade size is only weakly correlated with informedness, and that this correlation has declined materially over time. The accuracy of classifying trades below \$100,000 as uninformed falls below 50% after 2018. Even above \$100,000, misclassification remains substantial, averaging 23% over the sample period. Moreover, at round-number trade sizes (\$25,000, \$50,000, and \$75,000), the distributions of informed and uninformed trades are statistically indistinguishable. If the goal is to proxy for investor informedness, our findings caution against the continued use of the \$100,000 trade-size threshold in both academic research and regulatory analysis.

Our paper contributes to a large literature on market transparency and transaction costs in the municipal bond market.¹ Beginning with [Green et al. \(2007\)](#) and [Schultz \(2012\)](#), it is well established that municipal bonds trade in a decentralized broker–dealer market characterized by costly price discovery and heterogeneous investor sophistication. While prior studies document elevated and dispersed transaction costs (e.g., [Harris and Piwowar, 2006](#); [Green et al., 2006](#); [Green et al., 2007](#)), the evidence on long-run trends in new-issue markups has been mixed (see [Schultz, 2012](#); [Chalmers et al., 2021](#); [Wu, 2018](#); [Griffin et al., 2023](#)). We advance this literature by documenting a substantial decline in new-issue markups and by estimating, from a structural model of the rent mechanism, the two margins that

¹See [Bergstresser \(2022\)](#) and [Cestau et al. \(2019\)](#) for comprehensive surveys.

generate underwriters’ trading rents: the prevalence of uninformed trading and the markups the uninformed pay.

We also contribute to the literature on underwriting frictions. [Luby and Moldogaziev \(2013\)](#) link underwriting fees to market risk and show that local financial advisors mitigate these frictions. [Garrett \(2024\)](#) documents the effects of conflicts arising from dual underwriter–advisor arrangements. [Garrett et al. \(2023\)](#) shows that underwriters extract fewer information rents as the tax advantage of tax-exempt bonds increases. [Green et al. \(2007\)](#) highlights the substitutability between underwriting spreads charged to issuers and markups charged to investors. Our analysis deepens this strand by jointly examining both sources of compensation and showing that underwriters face increasing pressure on both sides: declining markups are not offset by rising underwriting spreads.

Finally, our results relate to a growing literature on clientele effects in municipal bond markets. [Babina et al. \(2021\)](#) and [Bergstresser and Cohen \(2015\)](#) highlight how tax incentives drive investor segmentation and retail concentration, while [Cornaggia et al. \(2022\)](#) show that inattentive investors contribute to persistent markup dispersion. We add to this literature by providing a new empirical measure of investor informedness and demonstrating that uninformed trading has declined markedly over the past two decades. Moreover, we show that the commonly used \$100,000 trade-size cutoff performs poorly as a proxy for investor informedness, particularly in recent years.

The paper proceeds as follows. Section 2 lays out the [Green \(2007\)](#) model, defines the equilibrium objects we measure, presents the structural model that estimates them, and delineates the paper’s contribution. Section 3 describes the data and documents the reduced-form declines in markups and underwriting spreads. Section 4 presents the empirical results.

Section 5.3 applies the estimated model to assess the \$100,000 trade size cutoff in classifying informed and uninformed trades. Section 6 concludes.

2 Conceptual Framework and Structural Model

2.1 Conceptual Framework

To understand the sources and dynamics of underwriter compensation, we build on [Green \(2007\)](#), who models the strategic interactions among issuers, underwriters, and investors. In the model, two intermediaries—each functioning as both underwriter and dealer—compete in a Bertrand manner (i.e., on price) for an issuer’s bonds and subsequently resell them to a mix of retail and institutional investors. In the resale market, dealers are free to set prices and discriminate between investor types. Retail investors have relatively high valuations but must pay a cost to learn dealers’ prices, a cost [Green \(2007\)](#) interprets as a measure of market transparency (public posting of quotations or low-cost access to transaction prices lowers it). Institutional investors face no such cost and have infinitely elastic demand at a lower price. Each underwriter–dealer serves the most profitable retail trades first and sells any remaining inventory to institutions, but its retail reach is limited by its distribution network. This capacity constraint disciplines primary-market bidding: acquiring more bonds than the retail network can absorb forces the excess to be sold at the lower institutional price, so underwriters avoid bidding too aggressively for the issuance.

The model admits a range of Nash equilibria, from a unique zero-profit equilibrium to the most profitable one, in which both dealers bid the institutional price and fully utilize

their retail capacity. [Green \(2007\)](#) argues that the most profitable equilibrium is particularly natural for the muni market setting, consistent with empirical evidence ([Green et al., 2007](#)). In this equilibrium, dealer d 's expected profit is:

$$\pi(b_d, Q_d) = (\bar{p} - v)q\mu + (v - b_d)Q_d, \quad (1)$$

where b_d and Q_d are the bid price the dealer offers to the issuer and the allocated quantity of issuance, respectively, μ is the mass of retail customers, q is the probability that a retail customer chooses not to pay the search costs and remains uninformed, $\bar{p}$ is the reservation price above which no retail investor will buy the security, and v is the institutional price.

Equation (1) decomposes underwriter compensation into two components with direct empirical counterparts. [Green et al. \(2007\)](#) document that the price offered by underwriters to issuers, plus the negotiated underwriting spread, is highly correlated with the institutional price, so $(v - b_d)$ maps to the underwriting spread and $(\bar{p} - v)$ to the markup per bond in the new issue market. The second term of Equation (1) is therefore the observable, issuer-paid underwriting spread component. The first term, $(\bar{p} - v)q\mu$, is the retail-distribution rent earned in the new issue market: the markups collected from distributing newly issued bonds to investors who face costly price search. Because we do not observe underwriting costs such as overhead, inventory risk, or regulatory compliance, the terms “compensation” and “revenues” refer throughout to these two gross components.

We take this model to the data: the paper estimates the equilibrium objects in Equation (1) and characterizes their evolution over two decades. The challenge is that the rent term is not directly observable: investor type is latent, and the economically relevant distinction is informed versus uninformed pricing behavior, not a trade-size label. Following [Green et al. \(2007\)](#), we therefore estimate a finite-mixture model that separates a tightly priced informed

regime from a diffuse uninformed regime using the full conditional distribution of new-issue markups. Section 2.2 presents the model; here we need only its two outputs for every trade: the probability $\hat{u}_{it}$ that trade i executes in the uninformed pricing regime, and the fitted markup schedules of the two regimes. From these two outputs we define the four objects that the paper measures. Let t index an estimation period, i the new-issue customer-buy trades in t , and j the issues sold in t .

The first object is the prevalence of uninformed order flow,

$$q_t^{trade} = \frac{1}{N_t} \sum_{i \in t} \hat{u}_{it}, \quad (2)$$

where N_t is the number of trades in period t .

The second object is the expected markup wedge for trade i ,

$$\hat{\Delta}_{it} = E \left[y_{it}^U - y_{it}^I \mid x_{it} \right], \quad (3)$$

the difference between the fitted uninformed and informed markup schedules (y^U and y^I) evaluated at the trade's characteristics x_{it} . Because the rent channel is a surplus, the empirical implementation truncates the wedge at zero,

$$\hat{\Delta}_{it}^+ = \max\{\hat{\Delta}_{it}, 0\}, \quad (4)$$

which is the posterior-weighted analogue of the money left on the table in [Green et al. \(2007\)](#).

The conditional markup wedge and the rent below are computed with $\hat{\Delta}_{it}^+$.

The third object is effective uninformed volume,

$$m_t^U = \frac{\sum_{i \in t} \hat{u}_{it} TradeSize_{it}}{\sum_{j \in t} Issuance_{jt}}, \quad (5)$$

the share of issuance distributed through uninformed order flow, where $TradeSize_{it}$ is the size of trade i and $Issuance_{jt}$ the issuance amount of issue j . This object is the empirical counterpart of the product $q\mu$: the model does not separately identify the probability that a retail customer remains uninformed (q) from the mass of reachable retail demand (μ), and

for the economics of underwriter compensation only their product matters.

The fourth object is the conditional markup wedge on effective uninformed volume,

$$\Delta_t^U = \frac{\sum_{i \in t} \hat{u}_{it} \hat{\Delta}_{it} TradeSize_{it}}{\sum_{i \in t} \hat{u}_{it} TradeSize_{it}}, \quad (6)$$

the average markup wedge per dollar of uninformed volume.

These objects aggregate into the market-wide retail-distribution rent, expressed in bps of issuance:

$$\mathcal{R}_t = \frac{\sum_{i \in t} \hat{u}_{it} \hat{\Delta}_{it} TradeSize_{it}}{\sum_{j \in t} Issuance_{jt}} = m_t^U \Delta_t^U, \quad (7)$$

an exact accounting identity and the empirical counterpart of $(\bar{p} - v)q\mu$. The issuer-paid component S_t (underwriting fees) completes the decomposition: total observed compensation is $C_t = \mathcal{R}_t + S_t$, and the distribution-rent share is $Share_t = \mathcal{R}_t/C_t$. Equation (2) thus measures the prevalence of the uninformed pricing regime, Equations (5) and (6) its effective quantity and unit rent, and Equation (7) the resulting market-wide rent.

The model predicts how each of these objects responds to forces that have reshaped the market over our sample period. Rising transparency (post-trade dissemination, markup disclosure, and related initiatives; Wu, 2018) lowers the cost of becoming informed, so fewer customers remain uninformed. The institutionalization of the investor base (the migration of direct retail holdings into mutual funds, ETFs, and separately managed accounts; Cestau et al., 2019) shrinks the retail mass μ .² These forces should reduce the prevalence of uninformed trading q_t^{trade} and compress effective uninformed volume m_t^U . Transparency should also narrow the wedge $(\bar{p} - v)$, and hence the conditional markup wedge Δ_t^U , because when search is a credible threat, dealers must moderate the prices they charge even the customers

²Healy, Patricia, “The Rise of Separately Managed Accounts – 2021 Update,” November 2, 2021, available at <https://www.cumber.com/market-commentary/rise-separately-managed-accounts-2021-update>

who remain uninformed. The distribution rent $\mathcal{R}_t = m_t^U \Delta_t^U$ should therefore contract on both margins. Finally, the model says who gains and who loses from transparency. Bertrand competition implies that the winning bid remains tied to the institutional valuation v and keeps underwriters from recouping the loss through higher spreads. The model thus predicts a compressed distribution rent $\mathcal{R}_t$ with no offsetting movement in the underwriting spread S_t : rents lost in new-issue distribution accrue to investors at the expense of underwriters, and could ultimately harm issuers if underwriters withdraw from the market.

2.2 Structural Model

The objects defined above require separating two latent pricing regimes in the data. The need for a mixture model starts from a striking market fact documented by [Green et al. \(2007\)](#): for the same newly issued bond, at nearly the same time, some customer purchases execute at or near the reoffering price while others pay widely dispersed and materially higher prices. A single regression with a trade-size slope cannot reproduce this simultaneous dispersion; it is naturally represented as a mixture of a tightly priced informed regime and a diffuse uninformed regime. We therefore estimate a latent class model—a finite mixture of regression models—in which a classification equation determines the regime and regime-specific outcome equations describe markups.

Our empirical specification follows the economic foundation of [Green et al. \(2007\)](#). An investor decides to become informed if the benefit of being informed, e.g., better terms of trade, exceeds the cost of information acquisition. Denote by z_i^* the net benefit of being informed, which is latent but modeled as a function of observable bond and trade character-

istics ω_i , i.e., $z_i^* = \omega_i\delta + u_i$. The markup each type (U for uninformed and I for informed) pays on trade i is also a linear function of observable characteristics x_i :

$$y_i^U = x_i\beta^U + \epsilon_i^U, \quad (8)$$

$$y_i^I = x_i\beta^I + \epsilon_i^I. \quad (9)$$

The observed markup y_i for trade i is thus:

$$y_i = \begin{cases} y_i^U & \text{if } z_i^* < 0 \\ y_i^I & \text{if } z_i^* \geq 0 \end{cases} \quad (10)$$

Following [Green et al. \(2007\)](#), the conditioning variables x_i in the markup regressions include observable bond characteristics as well as trade characteristics such as trade size, the days since the offering, and the percentage of the issue that has been sold prior to the trade. The set of conditioning variables ω_i in the classification regression includes all of the above plus a dummy for whether the issuer is a first-time issuer, the number of issues sold prior to this bond sale by the issuer, and the underwriting spread. The first two additions capture the information asymmetry of the issue, an important consideration in the decision to become informed. We also include the *National Underwriter* dummy in both parts of the model to account for the potential that national underwriters have different investor clientele and markup/markdown strategies.

The role of underwriting spread deserves some comment. The spread is agreed at issue pricing and is therefore predetermined relative to the customer trades the model classifies, so realized markups cannot reverse-cause it. It is nevertheless an equilibrium variable: issuers and underwriters set it while anticipating distribution capacity, expected clientele, and rent opportunities in the new issue market. Following [Green et al. \(2007\)](#), we include the spread in the classification equation to capture this equilibrium sorting—underwriters that

are compensated through a higher transparent spread are less likely to rely on uninformed distribution rents. We do not interpret its coefficient as the causal effect of spreads on informedness.

The first output that Section 2.1 requires is the trade-level uninformed probability, which the fitted model delivers as a posterior. Let $\pi_i = \Lambda(\omega_i \hat{\delta})$ denote the fitted prior probability that trade i is informed (Λ is the logistic CDF) and $f_I(y_i | x_i)$ and $f_U(y_i | x_i)$ the fitted conditional markup densities of the two regimes. The posterior informed probability is

$$\tau_i = \frac{\pi_i f_I(y_i | x_i)}{\pi_i f_I(y_i | x_i) + (1 - \pi_i) f_U(y_i | x_i)}, \quad (11)$$

and $\hat{u}_i = 1 - \tau_i$ is the model-implied uninformed probability that enters Equations (2)–(7).

The posterior log odds admit an exact decomposition:

$$\log\left(\frac{\tau_i}{1 - \tau_i}\right) = \log\left(\frac{\pi_i}{1 - \pi_i}\right) + \log\left(\frac{f_I(y_i | x_i)}{f_U(y_i | x_i)}\right), \quad (12)$$

in which the first term is the odds that a trade is informed based on its observable characteristics alone, and the second updates those odds according to which regime's fitted density better explains the markup actually paid. The posterior is therefore not a repackaging of observable covariates, as it also incorporates direct evidence on where the realized markup falls in the two conditional distributions. The second output of the fitted model is the regime schedules $x_i \hat{\beta}^U$ and $x_i \hat{\beta}^I$ whose difference gives the wedge $\hat{\Delta}_i = x_i(\hat{\beta}^U - \hat{\beta}^I)$. All important constructs described in Section 2.1 follow.

It is worth emphasizing that identification in our model comes from the full conditional distribution of markups: the two regimes are distinguished by their fitted dispersions, with the lower-variance regime labeled informed as in [Green et al. \(2007\)](#). This is consistent with the intuition that informed customers—institutions, advised accounts, or households that search—transact on comparatively homogeneous terms near the reoffering price, whereas

observationally similar uninformed customers pay a diffuse range of markups. This identifying structure holds strongly in the data: the dispersion ordering $\hat{\sigma}_I < \hat{\sigma}_U$ obtains in every estimation period, the two investor types are sharply separated, and nearly all trades receive posterior probabilities close to zero or one, so the model classifies almost every trade decisively (see Online Appendix (OA) Table A.1 for the full diagnostics).

A natural alternative to the mixture model is a single regression of markups on characteristics and an investor-type proxy. Such a regression, however, estimates only the probability-weighted average of the two regimes, $\pi_i E[y_i | I, x_i] + (1 - \pi_i) E[y_i | U, x_i]$. The average alone cannot reveal whether a decline in markups comes from underwriters encountering uninformed customers less often, from smaller markups charged to the customers who remain uninformed, or from both. Separating the two requires the regime probabilities and the regime-specific markup schedules, which our structural model can deliver.

Our contribution is threefold. First, we extend the measurement technology of [Green et al. \(2007\)](#) from a selected sample of about 1,000 deals issued in 2000–2003 (before real-time trade reporting) to the broad modern new issue market. Second, repeated estimation over rolling windows converts their cross-sectional decomposition into a two-decade time series of the extensive margin (m_t^U) and the intensive margin (Δ_t^U) of retail-distribution rents. Third, by combining these estimates with a uniquely comprehensive underwriting-spread database, we evaluate both components of underwriter compensation expressed in Equation (1) jointly and provide a market-wide accounting of how an opaque rent channel has evolved.

3 Data and Descriptive Evidence

3.1 Data

Municipal bond issuance information comes from the Mergent Municipal Fixed Income Securities database (“Mergent”): issue size, offering date, issuer name, offering type (competitive versus negotiated), and the agents involved in the sale. Each issuance typically comprises multiple CUSIPs (“bonds”) at various maturities, and for each bond Mergent reports the issuance amount, coupon, maturity, bond type (GO versus revenue), use of funds, tax status, callability, credit rating, and insurance status.

Trade data come from the Municipal Securities Rulemaking Board (“MSRB”): timestamp, trade size, trade type (customer buy, customer sell, or interdealer), and price. Because we measure the markups underwriters and affiliated dealers charge when distributing new issues, we focus on customer buy trades until the bond issuance is fully absorbed or within 14 days of the offering date, whichever comes first—similar to [Schultz \(2012\)](#) and [Griffin et al. \(2023\)](#). Over 75% of bonds are fully sold within 14 days (OA Table [A.2](#)); trades beyond that window are excluded because their markups are likely contaminated by intervening interest rate movements, the same reason we do not adopt the 60-day window of [Green et al. \(2007\)](#). On average, the new issue market lasts roughly 3.5 days, with a median of 1 day and an interquartile range of 1 to 6 days.

Following the literature, the new-issue markup is the customer purchase price minus the offering price, expressed as a percentage of the offering price (a markup of 0.2 corresponds to 20 bps). Bond-level analyses use the trade-size-weighted average markup across a bond’s

new-issue trades.

Underwriting spreads come from the SDC Platinum database (“SDC”), supplemented with Bloomberg data for the many issues SDC does not cover. SDC reports gross spreads per \$1,000 of par, which we convert to a percentage of issue size for consistency with Bloomberg; spreads are recorded at the issue level. Although spread data are not available for every issue—despite typically being disclosed in the Official Statement—our combined dataset is, to our knowledge, the most comprehensive collection of underwriting spreads used in a study of municipal underwriter compensation.

Table 1 details our cleaning procedures and filters. The sample runs from 2005 (when MSRB trade data become available) through 2024 and consists of 190,225 issues and 2,114,399 bonds; markup analyses use 11,551,952 new-issue customer buy trades, and spread analyses use the 138,540 issues with available spread data.

3.2 Types of Underwriters

Among the agents involved in bond issuance, we focus on underwriters, and we distinguish them by type because the competitive environment assumed in Section 2 fits some underwriters better than others. The rent mechanism itself requires only dealer pricing power over retail investors during distribution and heterogeneous information acquisition costs, which should apply across all types. The competition each type faces, however, differs: the Bertrand competition in Green (2007) is most applicable to national underwriters, which compete against close substitutes across overlapping markets; localized underwriters may instead compete on relationships, local retail access, or issuer-specific information. The model

should deliver sharpest predictions for the national segment, which accounts for roughly half of issuance in the data. We therefore expect heterogeneity across underwriter types and examine our results by type throughout, especially national underwriters versus the rest.

We classify underwriters into types by the geographic reach of their activity, i.e., the number of states in which they underwrite bonds in a given year. This is a natural dividing line: the municipal market is sharply segmented along state lines (Babina et al., 2021), and dealer-network evidence shows that the segmentation extends to intermediaries, with geographically specialized dealers concentrated in a few adjacent states (Jotikasthira et al., 2025).³ Geographic coverage thus proxies for an underwriter’s operational scale, capital base, and network reach. After manually cleaning underwriter names to consolidate name variants, we sort underwriters into four categories: 1) single-state underwriters, active in only one state; 2) small regionals, active in 2–19 states; 3) large regionals, active in 20–39 states; and 4) nationals, active in 40 or more states.

3.3 Descriptive Statistics

Table 2 summarizes underwriting activity, markups, and spreads by underwriter type over 2005–2024. Panel A shows the vastly different scales of operation: in an average year, a single-state underwriter does about seven deals worth \$64 million, a small regional about 52 deals (\$619 million), a large regional about 253 deals (\$5.6 billion), and a national underwriter

³Our classification is conceptually similar to the central-versus-peripheral dealer distinction of Li and Schürhoff (2019), who document a centrality premium in secondary-market markups. We cannot construct network-based centrality measures because the MSRB data version that includes primary market trades omits dealer identifiers; see Jotikasthira et al., 2025 for dealer-identified evidence on secondary-market specialization.

about 656 deals worth roughly \$19 billion across 44 states.⁴

Collectively, about eight national underwriters account for nearly 50% of annual volume and about 17 large regionals for another 30%, with volume shares exceeding deal-count shares because both groups concentrate in very large sales; the numerous small regional and single-state underwriters (95 and 113, respectively) underwrite much smaller issues and collectively account for only 20%. Figure 1 plots the time series: amid generally rising issuance, the national share climbs from below 40% in 2005 to about 60% in 2020 before falling back to about 50% in 2024, with the losses in 2018 and 2022–2023 absorbed mainly by large regionals.

Panel B of Table 2 turns to markups. Most bonds are distributed at the offering price—the median markup is zero for every underwriter type—but many trades carry positive markups, which are lowest and least variable for national underwriters. More importantly, Panel A of Figure 2 shows that except for single-state underwriters, average markups steadily decline, from the high teens (in bps) in 2005 to about five bps in 2024 for national and large regional underwriters. The decline could reflect more bonds distributed at zero markup, lower markups on marked-up bonds, or both—a distinction we explore below.

Panel C shows that national underwriters also charge the lowest underwriting spreads—about 75 bps of issue size, versus 90, 112, and 161 bps for large regional, small regional, and single-state underwriters, respectively. Because nationals underwrite much larger deals, their dollar fees per sale remain the highest even as fixed costs are spread over more volume. Panel B of Figure 2 shows the spreads of national and large regional underwriters holding in

⁴OA Figures A.1, A.2, and A.3 document the composition of underwriters by state in terms of the number of underwriters, quantity of deals, and volume, respectively.

the 50–100 bps range with only a slight recent increase, while those of smaller underwriters are consistently higher and rise faster.

Finally, we compute total underwriting compensation per issue—markups charged to investors plus fees charged to the issuer, as a percentage of issue size. Figure 3 shows a clear decline, especially after the 2008 financial crisis and among national underwriters, and most pronounced for the two eventual exiters, UBS and Citigroup, suggesting they were under significant pressure.⁵ In unreported time-trend tests, we confirm that the negative trend in Figure 3 is statistically significant, and more so for national underwriters. Overall, the descriptive evidence shows declining underwriter compensation, with underwriting fees not growing enough to offset falling markups.

3.4 Markups and Spreads After Controls

The declines documented above could in part reflect changing bond, issuer, and macro characteristics rather than changes in pricing. We therefore estimate conditional-mean regressions of the two compensation components—markups (at the bond level) and underwriting spreads (at the issue level, reflecting compensation collected from issuers rather than investors)—on year indicators and a full set of controls. The exercise serves a second purpose: the year effects indicate how underwriter compensation changes over time with composition held fixed, but, as Section ?? points out, conditional mean markups cannot separate how much of the evolution is due to changes in the probability of uninformed trades from changes in the markups on those trades. The regressions thus set the benchmark that the structural

⁵The sharp decline in 2024 reflects a surge in issuance—particularly larger deals—which mechanically inflates the denominator when compensation is expressed as a percentage of issue size.

estimates in Section 4 are designed to improve upon.

The regression model for the markup on bond i issued at time t is:

$$\begin{aligned} \text{Markup}_{i,t} = & \delta' \text{Year}_t + \beta_1 \text{National}_i + \beta_2' \text{National}_i \times \text{Year}_t + \beta_3 \text{LnIssueSize}_i \\ & + \beta_4 \text{LnBondSize}_i + \beta_5 \text{Maturity}_i + \beta_6 \text{Maturity}_i^2 + \beta_7 \text{Coupon}_i \\ & + \beta_8 \text{Premium}_i + \beta_9 \text{Callable}_i + \beta_{10} \text{Insurance}_i + \beta_{11} \text{AddCredit}_i \\ & + \beta_{12} \text{BankQlf}_i + \beta_{13} \text{OfferType}_i + \gamma_1' \text{BondType}_i + \gamma_2' \text{Rating}_i \\ & + \gamma_3' \text{Proceeds}_i + \eta' \text{State}_i + \epsilon_{i,t}, \end{aligned} \quad (13)$$

where $\text{Markup}_{i,t}$ is the trade-size-weighted average markup on new-issue customer buy trades as defined above, δ is a vector of year fixed effects capturing average markups by year after controls, and National indicates a national underwriter, included to test whether their markups differ once bond characteristics are held fixed.

Controls include logged issue and bond size (LnIssueSize , LnBondSize); Maturity and its square (to capture non-linearity); the Coupon rate; dummies for premium pricing (Premium), callability (Callable), insurance (Insured), additional credit enhancement (AddCredit), bank qualification (BankQlf), and negotiated offering (OfferType); and vectors (γ) of bond-type (GO, Revenue, or Other), rating, and use-of-proceeds indicators (General, Higher Education, Primary Education, Water, and Other), plus state fixed effects.

We plot the year fixed effects (δ_t) in Panel A of Figure 4. Each coefficient is the average new issue markup relative to the 2005 base year. Markups rise after the 2008 financial crisis through 2015 and then decline steadily to near their 2005 level.⁶ Once bond characteristics are held fixed, markups thus do not fall as early as the raw patterns in Figure 2 suggest, nor

⁶Two comments on these estimates are in order. First, the near-zero-rate period of 2010–2021 does not mechanically depress markups: they are measured as a percentage of the offering price rather than in yield terms, the regressions control for coupon, maturity, and premium status, and the decline is, if anything, steeper after 2018 as rates normalized. Second, trending bond characteristics do not drive the year effects: the only notable compositional shift is the post-2008 collapse of bond insurance (Cornaggia et al., 2023), a one-time break that cannot explain the persistent downward trend.

soon after the public dissemination of municipal trade data in 2005, the first of several major transparency initiatives. Panel B of Figure 4 plots β_2 , the national-underwriter markup differential. In most years, national underwriters' markups are significantly lower than other types', consistent with the unconditional statistics, even after the full set of controls.

The regression model for underwriting spreads on bond issue j issued at time t is:

$$\begin{aligned} \text{UndSpread}_{j,t} = & \delta' \text{Year}_t + \beta_1 \text{National}_j + \beta_2' \text{National}_j \times \text{Year}_t + \beta_3 \text{LnIssueSize}_j \\ & + \beta_4 \text{Callable}_j + \beta_5 \text{Insurance}_j + \beta_6 \text{AddCredit}_j \\ & + \beta_7 \text{BankQlf}_j + \beta_8 \text{OfferType}_j + \gamma_1' \text{BondType}_j + \gamma_2' \text{Rating}_j \\ & + \gamma_3' \text{Proceeds}_j + \eta' \text{State}_j + u_{j,t}, \end{aligned} \quad (14)$$

where the controls are as in Equation (13), excluding bond-level variables since the spread is an issue-level outcome.

Figure 5 plots the estimated year fixed effects and their National interactions for underwriting spreads. Panel A shows spreads moving somewhat opposite to markups—declining in the post-2008 period (as markups rose) and rising from 2016 onward (as markups fell)—suggesting a preference for stable total compensation, with shortfalls on one side partly offset on the other. However, Panel B shows that national underwriters' spreads are significantly lower than other underwriters' and trend down over time, mirroring their markup differential in Panel B of Figure 4. Collectively, these plots paint a consistent picture of declining total compensation, especially for national underwriters. What the year effects cannot reveal is what lies behind the decline. The structural estimates in Section 4 address this by measuring its underlying margins.

4 Empirical Results

This section presents our main empirical results based on the structural model of Section 2.2. To trace the evolution of the equilibrium objects discussed there, we estimate the model separately in overlapping two-year windows from 2005–2006 through 2023–2024. We use the expectation-maximization (EM) algorithm to obtain starting values and then estimate the full system by maximum likelihood. The estimation uses all customer buy trades in the new issue market for investment-grade bonds, because ratings enter both the class-membership and outcome equations.⁷ OA Tables A.3–A.4 report the estimated classification and uninformed-regime markup coefficients for some representative windows.

Below, we first present the evolution of the two estimated margins of the new issue market rent: the prevalence of the uninformed pricing regime and the markup earned in that regime. Second, we combine these estimates with data on trade sizes and underwriting spreads to decompose underwriter compensation and quantify the share attributable to retail-distribution rents. Third, we study the equilibrium relation between issuer-paid spreads and the pricing regimes. Next, we examine how distribution rents vary across underwriter type over time, since differences in retail distribution reach map directly into the rent mechanism of Green (2007). Finally, we present a battery of robustness checks addressing potential concerns with the main analysis.

⁷Below-investment-grade securities constitute only 0.25% of rated CUSIP-years in the structural input sample and never exceed 0.46% in any year, so this restriction avoids sparse rating categories without materially narrowing market coverage. Estimating the model separately for GO and revenue bonds yields similar results.

4.1 Time Series Results

This section reports how the prevalence of the uninformed pricing regime and the markup charged in that regime have evolved time. The prevalence of uninformed trading can fall due to either retail investors becoming genuinely more informed or retail capital being delegated to professional managers (e.g., separately managed accounts) such that trading is effectively conducted by institutional desks. Our estimates capture both, because the model classifies trades by pricing behavior rather than ownership, but either way, it is order flow on which the underwriter cannot earn the uninformed markup. Thus, for measuring underwriter compensation, the distinction is not consequential.

Figure 6 plots the prevalence series, computed over all transactions in each two-year estimation window and dated by the second year of the window, with 95% confidence intervals based on the bond-clustered covariance matrix. The model-implied incidence of uninformed trading falls by approximately 55% over the modern reporting era, from 57.5% in 2005–2006 to 25.7% in 2023–2024. The decline is not confined to a single regulatory episode or market cycle. Despite modest reversals in several estimation windows, the series exhibits a pronounced secular contraction.⁸ This is the direction the comparative statics of Section 2 predict: lower costs of becoming informed and a smaller retail mass both leave fewer customers uninformed.

Figure 7 plots the markup charged in the uninformed regime, based on the fitted value of the uninformed-regime outcome in Equation (8) averaged over all transactions in each

⁸OA Figure A.4 plots the probability of trading with uninformed investors for revenue bonds, which are riskier than GO bonds because they are backed by project-specific revenues rather than general taxing power. The results are analogous: the probability of uninformed trading falls from 59% in 2005–2006 to 33% in 2021–2022.

estimation window and expressed in bps. The series rises from approximately 102 bps in 2005–2006 to a crisis-period peak of 152 bps in 2008–2009, then falls to 79 bps in 2023–2024, a contraction of approximately 48% from the peak.⁹ This decline is likewise consistent with the comparative statics of Section 2: the more credible threat of search disciplines what dealers can charge even those investors who remain uninformed.

4.2 A Decomposition of Underwriter Compensation

We now translate the statistically estimated pricing regimes into the economically weighted components of underwriter compensation in Equation (1): effective uninformed volume, the conditional markup wedge on that volume, and the market-wide retail-distribution rent (Equations (5)–(7)). For this decomposition, we retain only trades issued in the second year of each fitted window, so that annual totals are non-overlapping. The issuer-paid component is measured as aggregate underwriting spreads per dollar of aggregate issuance, so that both compensation channels share the same market-wide denominator. Across all years, the sample used for the decomposition contains 8,465,941 customer buy trades in the new issue market for investment-grade issues, representing 109,373 issue-years.

Table 3 reports the decomposition by year. The extensive margin contracts sharply: effective uninformed order flow falls from 19.87% of aggregate issuance in 2006 to 7.36% in 2024, a 63.0% decline. This dollar-based contraction is even steeper than the decline in the count-based incidence of uninformed trading in Figure 6, reflecting the changing size composition of trades as well as the posterior probability assigned to each transaction.

⁹OA Figure A.5 plots the markup of uninformed trades for revenue bonds. These markups have declined considerably over time too.

The intensive margin also contracts materially. The conditional markup wedge declines from 80.47 bps in 2006 to 50.45 bps in 2024, a 37.3% reduction. From its 2008 peak of 118.87 bps, it falls by 57.6%. These proportional changes are at least as large as those implied by the average markups in Figure 7. The lower levels in the table are not evidence of a weaker effect: they arise because the decomposition places less weight on small transactions and on observations that would receive a high hypothetical uninformed-regime markup but have little posterior probability of actually belonging to that regime.

The product of the two margins produces a much larger contraction in the aggregate retail-distribution rent channel. The market-wide rent equals 15.99 bps of aggregate issuance in 2006, rises to 23.30 bps during the 2008 dislocation, and falls to 3.71 bps in 2024—a 2006–2024 decline of 76.8%. In dollar terms, underwriters earn approximately \$15.99 per \$10,000 of issuance through uninformed order flow in 2006, compared with \$3.71 in 2024. The result is stronger than a decline in average markups would indicate: the market generates substantially less price-discriminable order flow, and the order flow that remains uninformed yields a substantially smaller conditional rent.

The issuer-paid underwriting spread does not offset this erosion. The issuance-weighted underwriting spread declines from 55.80 bps in 2006 to 41.83 bps in 2024, so total observable compensation falls from 71.80 to 45.54 bps of aggregate issuance, and the distribution-rent share of underwriter compensation declines from 22.27% to 8.15%. The joint pattern shows who bears the loss. The rent is earned on the investor side of the underwriter’s business, and that is where it disappears: investors are less likely to be uninformed, those who remain uninformed pay smaller markups, and the saving accrues to investors. On the issuer side, competition for the issuer’s bonds pins down the bid and with it the spread (Section 2),

so underwriters cannot recoup lost distribution rents by charging issuers more, and the spread series confirms that they have not. The forces that make investors better informed thus benefit investors directly but weaken the economics of underwriting. Issuers are not insulated from that loss: if distribution no longer pays, underwriters may retrench or exit, which ultimately could negatively affect issuers in future bond sales.

The series above show both margins of underwriting rent contracting, but they cannot apportion the decline between them, because the rent is their product and the two margins interact. The rent identity in Equation (7) allows an exact answer: the change in the rent between two periods can be allocated to the two components, with their interaction split symmetrically between them (an exact two-factor Shapley allocation; [Shorrocks, 2013](#)). We use the pooled first four years (2006–2009) and last four years (2021–2024) to reduce sensitivity to any single year.

Table 4 reports the results of this quantification exercise. The distribution rent falls from 19.14 to 5.48 bps, a contraction of 13.66 bps, or 71.4%. Of this contraction, the decline in effective uninformed order flow from 19.05% to 8.70% accounts for -8.46 bps, or 61.9%, and the shrinking of the conditional markup wedge from 100.47 to 62.96 bps accounts for the remaining -5.21 bps, or 38.1%. Thus, roughly three-fifths of the lost rent reflects uninformed order flow that underwriters can no longer capture, and the remaining two-fifths reflects weaker pricing power over the order flow that stays uninformed. Neither contribution is negligible, and neither could be read off a declining average markup, so the erosion of the rent is not simply a relabeling of a single declining markup series.

A caveat to note is that our analysis captures only the two observable components of underwriter compensation. Underwriters may recoup some lost markup revenue through

less visible channels, e.g., advisory fees bundled into broader client relationships or favorable terms on future bond sales. But such channels are unlikely to fully offset the declines we document. The exits of UBS and Citigroup from municipal underwriting (two of the largest national underwriters, with presumably the most valuable issuer relationships) provide suggestive evidence that relationship-based income falls short of offsetting lost markup revenues. We therefore interpret the compression in observable compensation as a genuine reduction in underwriter revenues, while acknowledging that a complete accounting would require data on the full scope of underwriter–issuer relationships.

4.3 Underwriting Spreads and Equilibrium Sorting

The inclusion of underwriting spreads in the classification equation allows us to investigate whether spreads charged to the issuer and markups charged to uninformed investors are complements or substitutes. As discussed in Section 2.2, we read the coefficient as an equilibrium sorting relation rather than the causal effect of the spread. Figure 8 plots this coefficient over time, oriented as the log odds of the uninformed regime. The coefficient is significantly negative in every window, indicating that the more likely an underwriter expects to encounter uninformed investors and make profitable trades, the lower the underwriting fees it charges the issuer, allowing it to be competitive vis-à-vis other underwriters. This result extends the central cross-sectional finding of [Green et al. \(2007\)](#) to the modern new issue market.¹⁰

Over time, this substitution effect appears to lessen: the coefficient becomes less negative, from about -1.2 at the start of the sample to about -0.45 at the end, consistent with the

¹⁰In their switching equation, the underwriter discount raises the probability of informed trading; our coefficient is negative because the class-membership equation is oriented toward the uninformed regime.

cross-channel relation weakening as the rent channel contracts and competitive pressures prevent underwriting spreads from rising fast enough to compensate for lost markup revenues.

4.4 Heterogeneity by Underwriter Type

The market-wide contraction can mask different experiences across underwriter types. The estimated model already speaks to such heterogeneity: while national underwriters are no less likely than their counterparts to face uninformed investors at the beginning of our sample, they are between 0.10 and 0.40 percentage points less likely to trade with uninformed investors in recent years, and generally have significantly lower markups on uninformed trades than regional underwriters (see OA Figures A.6 and A.7). To characterize this heterogeneity in terms of the compensation objects, we recompute the full decomposition within each underwriter type using the estimated model’s posteriors and fitted markup schedules. Differences across types therefore reflect the composition and pricing of each segment’s order flow under a common structural map.

Figure 9 plots the annual distribution rent and underwriting spread by underwriter type, and OA Table A.5 reports the decompositions pooled over the first four years (2006–2009) and last four years (2021–2024). The erosion concentrates where the model predicts: for national underwriters, effective uninformed volume falls from 16.9% to 7.7% of issuance, the conditional markup wedge from 97.1 to 60.5 bps, and the rent from 16.38 to 4.63 bps, a 72% contraction. Large regional underwriters experience a nearly identical proportional decline in distribution rents (from 19.90 to 6.21 bps) and small regionals a substantial one (from 26.58 to 10.78 bps). Spreads do not offset the disappearing distribution rent for most

underwriter types. In particular, total compensation in the national segment falls from 70.2 to 45.9 bps of issuance.

The single-state segment is a striking exception: its rent is flat (26.58 versus 26.09 bps) and its spreads rise sharply (66.2 to 111.7 bps), so total compensation in this segment increases even as it collapses everywhere else. Because all estimates come from the same fitted pricing schedules, the contrast reflects the segments' order flow rather than differences in estimation. It indicates that the rent has eroded where distribution is national and institutional-facing, while a segment of locally distributed issues (small issuers served by underwriters active in a single state, whose natural clientele is local retail investors) continues to generate uninformed order flow at rates resembling the early sample. A caveat to note is that the single-state segment is small and selected (roughly seven deals and \$64 million per underwriter-year as reported in Section 3.3), and its annual series are volatile.

These results also contextualize the retrenchment of national underwriters documented in Section 3: the exits and share losses occurred precisely in the segment where the rent channel has come closest to disappearing.

4.5 Robustness

The results above could be sensitive to how trades are assigned to the two pricing regimes, to the changing mix of competitive and negotiated offerings, to bonds allocated at the reoffering price in negotiated sales, to the unwinding of the 2008–2009 dislocation, to how we identify trades in the new issue market, and to the weighting of issues by size. We examine each in turn using the observation-level predictions of the estimated models and summarize the

results here. Unless noted, early and late statistics pool 2006–2009 and 2021–2024 as in Table 4.

Regime classification. The results are not an artifact of how trades are assigned to the regimes. The baseline assigns each trade a posterior probability of belonging to the uninformed regime, which combines the trade’s class prior (from its characteristics, including the underwriting spread) with the fit of the two markup densities (Equation (12)). We compare three alternatives. The first gives every trade in a window the same class prior, the window mean, so that only the markup densities distinguish the regimes (we call this the density-only posterior). The second keeps the class prior but fixes every issue’s underwriting spread at its window mean, removing the spread’s role in the classification. The third keeps the baseline posterior, so both the characteristics and the markup densities still enter, but discards its fractional weights, so that each trade is assigned entirely to the regime with the higher posterior probability and thus no trade counts partly in both regimes (“hard classification”).

Panel A of OA Table A.6 shows that none of these changes the results: the annual mean probability of uninformed trading moves by at most 0.03 percentage points and the annual rent by at most 0.22 bps. The pooled early-to-late rent decline is essentially identical in every construction (e.g., 19.20 to 5.47 bps under the density-only posterior, versus 19.14 to 5.48 bps at baseline). The fit of the markup densities dominates the classification: cross-sectional variation in the class prior in general, and in the underwriting spread in particular, does not mechanically generate our time series. The agreement between this hard classification and the posterior-weighted baseline also shows that the results do not depend on how trades with intermediate posterior probabilities are weighted.

Offering methods and retail order period. The decline is not a composition effect of offering methods. Reconstructing the decomposition separately for competitive and negotiated offerings, the rent falls in both: from 17.12 to 9.01 bps (47%) among competitive sales and from 19.56 to 4.68 bps (76%) among negotiated sales (Panel B of OA Table A.6). That said, the source of the rent decline differs. Among competitive sales, the decline comes largely from the shrinking conditional markup wedge (55.07 to 35.63 bps), whereas among negotiated sales, it comes mostly from the probability of uninformed trading (54.5% to 15.6%).

One might argue that the retail order periods (ROP) of negotiated sales could produce the pattern of declining probability of uninformed trading. In a negotiated sale, the underwriter can fill retail orders at the reoffering price before final pricing. Such purchases carry no markup and are classified as informed, so a growing volume of them would lower the measured probability of uninformed trading for a reason that says more about distribution mechanics than about investors' informedness. To rule this out, we recompute the key equilibrium constructs using a subset of negotiated sales with little ROP activities. Because the MSRB transaction data do not flag these ROP trades, we compute the fraction of customer buy volume executed within 1 bps of the reoffering price through event day 1 and screen out issues where this fraction is at least 10%, 25%, or 50%. The screens are deliberately blunt (the tightest retains only 368 issues in the late period), and the retained issues carry higher rents than the market as a whole, so only the decline within them is informative. The pattern survives: the rent falls by 56.6%, 54.3%, and 40.2% under the three screens (Panel C of OA Table A.6). With the conditional markup wedge largely unchanged, it appears that the decline in rent is mostly due to a lower probability of uninformed trading, even among issues with little ROP volume. Thus, ROP allocations cannot account for our findings.

Post-crisis period. The contraction is not an artifact of the 2008–2009 dislocation unwinding. Measured from 2013, well after the crisis and its reversal, the rent still falls from 14.35 to 3.71 bps by 2024 (a 74% decline), effective uninformed volume from 12.90% to 7.36%, and the conditional markup wedge from 111.27 to 50.45 bps (Table 3).

Definition of new issue market trades. The results do not depend on how we define new issue market trades. The baseline counts customer buy trades until the issue is absorbed or through the fourteenth day after the offering, whichever comes first. A potential concern is that the later trades in that window are the ones most exposed to intervening interest rate movements and might contaminate markups. We therefore recompute the objects on narrower horizons, using the same posteriors and fitted schedules (OA Table A.7). The erosion appears within every horizon. Among trades on days 2–5 after the offering, the conditional markup wedge falls from 112.2 to 70.0 bps and the rent from 4.35 to 1.47 bps of aggregate issuance. Among trades on days 2–14, the corresponding declines are from 106.2 to 59.3 bps and from 10.44 to 3.18 bps. On the first five trading dates of each bond, the closest counterpart to the horizon in Green et al. (2007), the probability of uninformed trading falls from 51.9% to 22.5%, the wedge from 96.7 to 62.5 bps, and the rent from 17.00 to 5.23 bps. Observed raw markups and the hard-classified money left on the table decline similarly. We conclude that trades late in the fourteen-day window do not drive the results.

Weighting scheme. The erosion is not a by-product of weighting issues by size. As shown in OA Table A.8, weighting each issue equally instead, the posterior-weighted issue-mean rent falls from 21.70 bps in 2006 to 8.59 bps in 2024 while the issue-mean underwriting spread is flat (73.20 and 78.44 bps). The typical issue thus shows the same pattern as

the market-wide aggregates reported in Table 3: a collapsing rent, no offsetting movement in the spread, and a rent share of compensation falling from 19.9% to 8.6%. Issue-level means exceed the aggregates because large issues generate less rent per dollar and pay lower percentage spreads, but the trend is the same either way.

5 Applications of the Model

The estimated model is useful beyond decomposing underwriting compensation and identifying the sources of its decline. This section presents two applications: 1) testing the effects of transparency initiatives and increasing investor sophistication on the likelihood of being uninformed and the trade terms being offered on uninformed trades, and 2) evaluating the quality of the trade-size cutoff commonly used to classify investors.

5.1 The Effects of Transparency

One important trend in the municipal bond market is the increase in market transparency over time. Real-time transaction reporting commenced in January 2005 (the beginning of our sample), followed by continued efforts by the MSRB and SEC to improve transparency ever since. Although we cannot directly measure the degree of market transparency to establish a direct linkage between transparency and our documented evidence of increasing investor informedness, there are important rules implemented by the SEC and MSRB that significantly increase market transparency (Wu, 2018). Specifically, we exploit the implementation of amended MSRB Rules G-15 and G-30 and FINRA Rule 2232 in May 2018, which required dealers to disclose the markup in certain primary market transactions with

investors, as a shock to transparency.

It is worth noting that we did not use these transparency events to identify the secular decline in distribution rent documented above. Instead, they provide a market-wide comparative-static validation of the mechanism: transparency initiatives lower the cost of price information, so both markups and within-bond price dispersion (a measure of implied search costs for investors, as used in [Green et al., 2007](#)) should compress in response. In this section, we first verify the effects of the transparency shock on these two directly observed outcomes. We then ask which of the model’s primitives moves when the information environment improves.

We start by augmenting the regression model in Equation (13) for markups with a *Post* dummy variable reflecting the post-policy change, along with its interaction with the *National* dummy, omitting the *Year* $\times$ *National* fixed effects. We estimate a similar specification for price dispersion, which is computed as the standard deviation of trade prices in a given bond in the new issue market. Table 5 presents the results. Panel A shows that the transparency shock significantly reduces markups charged by national underwriters in the primary market, in line with the descriptive evidence in Figures 2 and 4 and consistent with [He \(2024\)](#), who finds that retail trading costs declined following the markup disclosure rule. Panel B shows a corresponding decline in price dispersion, indicating reduced search costs for investors. Together, these results suggest that greater transparency facilitates information acquisition, thereby limiting underwriters’ ability to charge high markups.

Next, to examine the effects of this transparency shock on the prevalence of uninformed trades and markups on those trades, we reestimate the same two-class finite-mixture model on a narrow window around the event. The pre-period is July 2017 through April 2018 and

the post-period is June 2018 through March 2019, both falling after the June 2017 issue-price regulations. May 2018 is omitted because the requirements took effect during that month. The specification otherwise follows, with two additions: $Post$ and $Post \times National$ enter the latent-class equation and both class-specific markup equations, and a common linear calendar-month trend enters both parts of the model to absorb gradual movement in the pricing environment. The model therefore does not force the response to appear as greater informedness. Increased transparency may change the probability of entering a specific pricing regime, either conditional markup schedule, both, or neither.

As before, from the estimates we compute four objects under the pre- and post-disclosure states, each evaluated over one common empirical distribution of bond, issue, transaction, and underwriter characteristics so that composition cannot drive the comparison. We use class priors rather than posteriors so that the transparency state is not conditioned mechanically on the realized markup being decomposed. Standard errors use the delta method and the bond-clustered covariance matrix from the joint maximum-likelihood estimate.

Panel A of Table 6 shows that the response is concentrated in the incidence of uninformed pricing. The probability of uninformed trading falls from 41.13% to 38.30%, a decline of 2.83 percentage points, or 6.9% of its pre-disclosure level. Effective uninformed volume falls from 14.84% to 12.86% of represented issue par, a decline of 1.97 percentage points, or 13.3%. Both changes are significant at the 1% level. The conditional wedge, by contrast, does not decline (it instead rises by an insignificant amount). The distribution rent thus falls from 15.22 to 13.72 basis points of issue par. The 1.50-basis-point decline is 9.9% of the pre-disclosure rent and is significant at the 5% level.¹¹

¹¹The informed component in the event-window mixture is estimated close to a mass point at the reof-

We decompose the 1.5 bps decline in rent into that due to changes in the uninformed order flow, evaluated at the average conditional markup wedge, and that due to changes in the conditional markup wedge, evaluated at the average uninformed order flow. The former contributes -2.07 bps to the decline, more than offset the contribution of $+0.57$ bps of the latter. This accounting is not an additional identification exercise, but it shows that reduced exposure to the uninformed regime accounts for more than the entire decline in rent, counteracting the small increase brought about by the conditional markup wedge of uninformed trades.

In Green’s language, the response operates through q and $q\mu$ rather than through an immediate compression of $\bar{p} - v$, which is the margin the theory ties most directly to a lower cost of price information. The distinction is economically meaningful. Greater transparency need not lower every component of dealer pricing power at the same horizon. A disclosure rule can make it easier for marginal customers to locate or recognize the relevant benchmark price, moving some transactions from the uninformed to the informed regime, while leaving the pricing schedule faced by the remaining nonsearchers largely unchanged. The event-window estimates display that pattern.

Panel B shows that the structural response is not specific to national underwriters. Setting the *National* indicator to zero or one for every observation while holding the remaining covariate distribution fixed, the probability of uninformed trading falls by 2.79 percentage

fering price, consistent with the sharp concentration of informed transactions at zero markup in the rolling estimates. We therefore emphasize the economic pattern across the structural objects rather than the significance of any single nonlinear contrast. The analytic counterfactual predictions reproduce Stata’s native finite-mixture predictions to within 3.3×10^{-13} . Central finite-difference checks produce a maximum scaled discrepancy of 2.3×10^{-7} for the smooth objects and 2.5×10^{-5} for the objects involving the positive-part markup wedge, and fewer than 0.008% of observations lie within 10^{-4} percentage points of the wedge’s zero boundary.

points for other underwriters and by 2.86 percentage points for national underwriters, and effective uninformed volume falls by roughly two percentage points for both groups. None of the national differentials in $P(U)$, effective uninformed volume, the conditional wedge, or the distribution rent is statistically distinguishable from zero. The structural information response is market-wide: disclosure changes the search environment that customers face across the municipal distribution system rather than producing a shift in latent customer type or conditional pricing that is unique to national underwriters.

We interpret Table 6 as a structural decomposition of the event-window pricing response and not as a separate source of causal identification. Its contribution is nonetheless substantive. The model is estimated jointly over a narrow, locally detrended event sample and permits the rule to affect every structural margin, and within that design the data place the response in the q -channel: fewer customer dollars remain exposed to uninformed pricing, which lowers the aggregate retail-distribution rent even though the wedge paid by the residual uninformed customers does not measurably decline. Taken together, the evidence in this subsection indicates that the rise in investor informedness reflects, at least in part, an intrinsic improvement in information environment as a result of the implementation of transparency initiatives

5.2 The Effects of Increased Investor Sophistication

Section 5.1 shows that a public reduction in the cost of price information moves order flow out of the uninformed pricing regime but leaves the price paid by the customers who remain there largely unchanged. Public disclosure is only one source of investor informedness. Investors

also differ in their own capacity to acquire and use price information, and that capacity varies across states and over time. In this subsection we ask whether externally measured investor sophistication shows up in the model’s primitives, and if so in which one: a lower probability of uninformed trading, less volume exposed to uninformed pricing, or a smaller markup on the uninformed trades that remain.

We use two measures that capture different sources of sophistication. Direct sophistication is the household’s own financial capability, which we proxy with the prevalence of financial-market participation among a state’s tax filers. Delegated sophistication is capability acquired through professional help, which we proxy with the local density of investment advisers. The two need not act on the same margin. Direct sophistication could work on either one: it can lead more households to become informed, or it can improve the terms obtained by those who do not. Delegated sophistication is most naturally a regime-switching channel, since advisers and managed accounts let retail capital trade on institutional terms (Green et al., 2007), and should lower $P(U)$ and m^U .

We measure direct sophistication using IRS Statistics of Income data, available through 2022.¹² For each state-year we average the shares of returns reporting taxable interest, dividends, and capital gains. Because these shares rise steeply with income, we compute them within adjusted gross income (AGI) classes and aggregate with fixed national AGI-class weights, so that the measure, expressed in percentage points, reflects participation at a given income level and not the state’s income mix.

We measure delegated sophistication as the number of SEC- and state-registered invest-

¹²Individual tax statistics are available at: <https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-return-form-1040-statistics>.

ment advisers active in a state per 100,000 residents.¹³ SEC-registered advisers, those with at least \$100 million under management or operating in more than 15 states, are taken from annual Form ADV filings and counted as active in every state in which they make a notice filing. Smaller advisers register with state securities authorities through the Investment Adviser Registration Depository, and we use the listed regulators and the approval or suspension date of each registration to determine the state-years in which they are active. Because the 2024 filings are incomplete, the series ends in 2023.

For state s and year t we estimate the following regression:

$$Y_{st} = \alpha_s + \lambda_t + \beta_1 \text{Direct}_{s,t-1} + \beta_2 \text{Delegated}_{s,t-1} + \Gamma' X_{st} + \varepsilon_{st}, \quad (15)$$

where Y_{st} is one of the four structural objects of Section 5.1, computed from the rolling estimates over all new-issue transactions in bonds issued in state s in year t and weighted by posterior class probabilities. The controls are log population, the lagged share of upper-income returns, and log aggregate municipal issuance. We require at least five issues per state-year, weight by issue count capped at the annual 95th percentile, and cluster standard errors by state. The sample contains 841 state-years from all 50 states over 2007–2023. State and year fixed effects confine the comparison to within-state changes net of national developments such as the transparency reforms, but the design is not causal.

Table 7 shows that direct sophistication loads on the conditional wedge and on nothing else. A one-percentage-point increase is associated with a wedge that is 5.37 basis points smaller (standard error 1.47), and a within-state interquartile increase of about 5.3 percentage points with a wedge almost 29 basis points smaller. This result is not driven by

¹³Data on SEC-registered investment advisers are available at: <https://www.sec.gov/data-research/sec-markets-data/information-about-registered-investment-advisers-exempt-reporting-advisers>. Data on state-registered advisers are available at: <https://adviserinfo.sec.gov/compilation>.

short-lived fluctuations in the IRS measure. In long differences between 2007–2010 and 2020–2023, the direct-sophistication coefficient is -4.47 with equal state weights and -6.20 with issue-count weights (OA Table A.9), bracketing the panel estimate, and the implied effect of an interquartile long-run increase is a wedge 8 to 11 basis points smaller.

We find no evidence that direct sophistication moves transactions across regimes: its coefficients for $P(U)$ (0.05, standard error 0.67) and m^U (0.52, standard error 0.42) are insignificant. These nulls are imprecise, however. The 95% confidence interval for the interquartile effect on $P(U)$ spans roughly -6.7 to $+7.2$ percentage points. The rent coefficient is negative but insignificant (-0.40 , standard error 0.45), because direct sophistication lowers the price component of the rent but not its quantity component.

Direct sophistication therefore appears to improve the terms of uninformed trades rather than reduce their number. This is the pattern that search theory predicts: a clientele that could readily check a quote against the reoffering price or recent trades obtains better terms even if it seldom does so. The option to become informed disciplines dealer pricing without being exercised, which is why the wedge can fall while the share of uninformed trades does not. That said, we cannot rule out the alternative that sophisticated households are partially informed in ways that the model’s two-class structure assigns to the uninformed regime.

The evidence on delegated sophistication is inconclusive. Its coefficient for $P(U)$ has the predicted sign (-1.60 percentage points per ten additional advisers per 100,000 residents) but a standard error of 1.68, and it does not significantly predict any other object. We do not conclude that delegation is unimportant: adviser counts can be a noisy proxy for the municipal assets managed on behalf of resident households.

Together with Section 5.1, these results indicate that the two components of the dis-

tribution rent respond to different forces. Public disclosure reduces how much order flow reaches the uninformed regime, whereas direct sophistication reduces the price paid by the customers who remain there. A conventional regression of markups could at best show that markups are lower, but could not show that they do so through different channels, because the observed markup is a composite of the incidence of uninformed trading and its price, which only a structural model like ours can separate.

5.3 Application: Evaluating the Trade Size Cutoff

In many studies regarding municipal markets, trade size is typically used as the device to delineate retail from institutional investors, mainly for highlighting the difference in the level of information each type possesses. The most common threshold for such classification is \$100,000. Analyzing transaction costs in the muni market, one of the earliest and most influential studies to employ the \$100,000 threshold to differentiate between retail and institutional trades is [Harris and Piwowar \(2006\)](#).¹⁴ The SEC also adopts this threshold in Rule 15c2-12, which considers investors with minimum denominations of \$100,000 as “sophisticated” (see [Ivanov et al., 2024](#)). Because our model classifies, at the trade level, whether the investor is informed or uninformed, we can assess if the trade size classification is appropriate when the goal is to differentiate on the level of investor’s informedness.¹⁵

Specifically, we examine how strongly trade size and the probability of the investor be-

¹⁴Other notable papers include [Hong and Warga \(2004\)](#), [Green et al. \(2006\)](#), [Schultz \(2012\)](#), [Cuny \(2017\)](#), and [Hendershott et al. \(2020\)](#).

¹⁵It is important to point out that in [Green \(2007\)](#)’s model, the relevant distinction is between being informed versus uninformed rather than between retail versus institutional. In the model, due to the choice of becoming informed, informed retail investors obtain the same trade terms as institutional investors and thus are indistinguishable from the latter when it comes to the observed markups.

ing uninformed are related. If trade size is a reasonable means to classify informed and uninformed trades, the correlation should be strongly negative and the rate of correct classification (i.e., the fraction of uninformed trades among trades below \$100,000 and the fraction of informed trades among trades above \$100,000) should be closer to 1 (see Table 8).

For each trade size category, we report the percentage of uninformed trades, which indicates the correct classification rate for trades under \$100,000 (column 1) but incorrect classification for trades above \$100,000 (column 4). Column 1 reveals that the percentage of trades below \$100,000 being correctly classified as uninformed trades is about 60% on average over the full sample period, but this rate clearly decreases over time, falling below 50% after 2018. In addition, column 4 shows a non-trivial proportion of trades above \$100,000 that are not informed, ranging from 38% at the beginning of the sample period to about 14% at the end of the sample period. Even though the quality of classification improves over time for trades above \$100,000, the deterioration in the correct classification rate for trades below \$100,000 suggests that the \$100,000 cutoff is increasingly inappropriate for separating uninformed and informed trades.

We also compute the point-biserial correlation between trade size and uninformed trade and report these correlations separately by year for trades below and above the cutoff (see columns 2 and 5 of Table 8). While always negative (as theory would predict), the correlation is weak: its magnitude is below 0.1 in every year for trades above \$100,000 and only slightly above 0.1 in most years for trades below \$100,000 (the full-sample values are -0.068 and -0.124 , respectively). Trade size is thus not highly correlated with the probability of investor informedness on either side of the cutoff.

Another metric to evaluate the appropriateness of the trade size cutoff in correctly sepa-

rating trades into uninformed versus informed is the total variation distance (TVD) between the informed and uninformed distributions. TVD is a symmetric measure bounded between 0 and 1 that can broadly be interpreted as the total probability mass needed to “shift” from one distribution to the other in order to equalize them (Gibbs and Su, 2007). A lower value indicates that the two distributions are similar, whereas a high value indicates that they are different. This measure, reported in columns 3 and 6 of Table 8, shows again that the distance between the uninformed and informed trade distributions is declining dramatically over time and is much lower below \$100,000 trade sizes than above. The informed and uninformed distributions are nearly twice as “close” to each other below \$100,000 compared to above. This finding indicates that it is increasingly harder to delineate uninformed from informed trades, especially among those under the cutoff.

We zoom in on trades below \$100,000 and plot the overlapping distributions of uninformed and informed trades for the entire sample period using bin sizes of \$5,000 in Panel A of Figure 10. While there is a predominance of uninformed trades, there is a very substantial percentage of informed trades. At some trade sizes (e.g., \$25,000 and \$50,000), informed trades have even higher probability than uninformed trades. Furthermore, we conduct multiple testing across quantiles of the informed and uninformed distributions as in Goldman and Kaplan (2018). Panel B of Figure 10 depicts the empirical CDFs for each distribution, and the quantiles where we can reject equality of the CDFs (the bolded black bar on the x-axis) at a 5% familywise error rate. Importantly, round trade sizes of \$25,000, \$50,000, and \$75,000 are areas of the distribution where we cannot reject equality. This finding suggests that a much more improved method of using trade size to proxy for uninformed trade would use only non-“round” trade amounts below a \$100,000 threshold. Even then, the results on

correlation and TVD indicate that there could still be substantial misclassification.

Overall, the results in this section demonstrate that the correlation between trade size and whether the investor is informed is weak and declines over time. For trades under \$100,000, the distributions of informed and uninformed trades are increasingly indistinguishable, with similar probabilities at “round” trade sizes. Our analysis indicates that, at least in the new issue market, the conventional \$100,000 trade-size threshold may not reliably distinguish informed from uninformed trades.

6 Conclusion

Our study provides a comprehensive assessment of the evolution of the municipal bond market from 2005 to 2024, a period marked by substantial shifts in transparency, investor composition, and underwriting practices. After compiling detailed data on municipal bonds and their characteristics, we document a sustained decline in both underwriting spreads and new-issue markups, particularly among national underwriters, and show that these trends persist after controlling for an extensive set of bond-, issue-, and issuer-level characteristics.

To understand the forces driving these changes, we estimate a structural model in the spirit of [Green \(2007\)](#) and [Green et al. \(2007\)](#), which links observed markups to bond and trade characteristics while allowing the markup schedule to differ across informed and uninformed investors. Our estimates reveal a sharp contraction in the prevalence of uninformed investors—from more than 57% in 2006 to 26% in 2024—and a substantial decline in the markups underwriters can charge them. Together, these margins compress the retail-distribution rent from 15.99 to 3.71 bps per dollar of issuance, a 76.8% decline, of which a

Shapley decomposition attributes 61.9% to shrinking effective uninformed volume and 38.1% to the narrowing conditional markup wedge. Underwriting spreads do not rise to offset the loss and the rent’s share of total underwriter compensation falls from 22.27% to 8.15%.

Our results highlight several advantages of the structural model over conventional regression approaches for understanding the dynamics of distribution rents. The rent’s components are latent-regime quantities that average markups and trade-size classifications cannot reliably recover. Our model instead uses the full conditional distribution of markups to separate the informed from the uninformed pricing regime, which lets us track the prevalence of uninformed order flow and the markup wedge on that flow separately over time. The resulting series show how a central over-the-counter intermediation mechanism has responded to transparency shocks and changes in investor sophistication. Transparency works on the extensive margin by reducing the prevalence of uninformed trading, whereas household financial capability works on the intensive margin by compressing dealers’ pricing power over those who remain uninformed.

Finally, our results caution against the use of the \$100,000 trade-size threshold for distinguishing retail from institutional trades. Trade size is only weakly correlated with informedness and has become increasingly unreliable over time: classification accuracy for trades below \$100,000 falls below 50% after 2018, and misclassification remains substantial even above the threshold. If the goal is to proxy for investor informedness, the traditional cutoff performs poorly and should be used with caution in both academic work and policy analysis.

Our findings advance the understanding of long-run changes in municipal market microstructure by estimating the two margins of an opaque rent channel and documenting that both have contracted sharply: investors are less likely to trade uninformed, and those who

do pay smaller markups. They also highlight a potential unintended consequence of greater transparency and investor sophistication: these developments benefit investors by reducing information frictions but may leave muni issuers worse off if lower underwriter compensation lead underwriters to exit or retrench, as predicted by [Green \(2007\)](#).

References

- Babina, T., Jotikasthira, C., Lundblad, C., and Ramadorai, T. (2021). Heterogeneous taxes and limited risk sharing: Evidence from municipal bonds. *Review of Financial Studies*, 34(1):509–568.
- Bergstresser, D. (2022). The municipal bond market. *Research Handbook on Financial Markets*.
- Bergstresser, D. and Cohen, R. (2015). Changing patterns in household ownership of municipal debt: Evidence from the 1989-2013 surveys of consumer finance. Working paper, Brandeis University.
- Cestau, D., Hollifield, B., Li, D., and Schurhoff, N. (2019). Municipal bond markets. *Annual Review of Financial Economics*, 11:65–84.
- Chalmers, J., Liu, Y. S., and Wang, Z. J. (2021). The difference a day makes: Timely disclosure and trading efficiency in the muni market. *Journal of Financial Economics*, 139(1):313–335.
- Cornaggia, K., Hund, J., and Nguyen, G. (2022). Investor attention and municipal bond returns. *Journal of Financial Markets*, 60:100738.
- Cornaggia, K., Hund, J., and Nguyen, G. (2023). The price of safety: The evolution of municipal bond insurance value. *Management Science*, forthcoming.
- Cuny, C. (2017). When knowledge is power: Evidence from the municipal bond market. *Journal of Accounting and Economics*, 64(1):277–297.

- Garrett, D. (2024). Conflicts of interest in municipal bond advising and underwriting. *Review of Financial Studies*, forthcoming.
- Garrett, D., Ordin, A., Roberts, J. W., and Serrato, J. (2023). Tax advantages and imperfect competition in auctions for municipal bonds. *Review of Economic Studies*, 90(2):815–851.
- Gibbs, A. L. and Su, F. E. (2007). On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435.
- Goldman, M. and Kaplan, D. M. (2018). Comparing distributions by multiple testing across quantiles or cdf values. *Journal of Econometrics*, 206(1):143–166.
- Green, R. C. (2007). Presidential address: Issuers, underwriter syndicates, and aftermarket transparency. *Journal of Finance*, 62(4):1529–1550.
- Green, R. C., Hollifield, B., and Schürhoff, N. (2006). Financial intermediation and the costs of trading in an opaque market. *The Review of Financial Studies*, 20(2):275–314.
- Green, R. C., Hollifield, B., and Schürhoff, N. (2007). Dealer intermediation and price behavior in the aftermarket for new bond issues. *Journal of Financial Economics*, 86:643–682.
- Griffin, J. M., Hirschey, N., and Kruger, S. (2023). Do municipal bond dealers given their customers “fair and reasonable” pricing? *Journal of Finance*, 78(2):887–934.
- Harris, L. and Piwowar, M. S. (2006). Secondary trading costs in the municipal bond market. *Journal of Finance*, 61:1361–1379.

- He, Y. (2024). Mark-up disclosure and trading costs in the municipal bond market. *Working paper, University of Nebraska–Lincoln*.
- Hendershott, T., Li, D., Livdan, D., and Schürhoff, N. (2020). Relationship trading in otc markets. *Journal of Finance*, 75(2):683–734.
- Hong, G. and Warga, A. D. (2004). Municipal marketability. *Journal of Fixed Income*, 14(2):86–95.
- Ivanov, I., Zimmermann, T., and Heinrich, N. W. (2024). Limits of disclosure regulation in the municipal bond market. *Management Science*, (forthcoming).
- Jotikasthira, C., Lundblad, C. T., and Xue, J. (2025). Dealer specialization and market segmentation. *Working paper*.
- Li, D. and Schürhoff, N. (2019). Dealer networks. *Journal of Finance*, 74(1):91–144.
- Luby, M. and Moldogaziev, T. (2013). An empirical examination of the determinants of municipal bond underwriting fees. *Municipal Finance Journal*, 34(2):13.
- Schultz, P. (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.
- Shorrocks, A. F. (2013). Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. *Journal of Economic Inequality*, 11(1):99–126.
- Wu, S. (2018). Transaction costs for customer trades in the municipal bond market: What is driving the decline? Technical report, Municipal Securities Rulemaking Board.

Table 1: Sample construction and data filters

This table shows the filters applied in constructing the bond sample and the bond-trade sample. Bond data are from the Mergent municipal fixed-income database. Trade data is from the MSRB. We only consider trades in the new issue market, defined as those within 14 days of the offering date or when the bond size is fully absorbed by customer buy trades, whichever comes first. The sample period is from 2005 to 2024.

Filters	# Issues	# CUSIPs	# Trades
Full Mergent bond universe (as of 2/21/2025)	524,644	4,710,200	
<i>Cleaned using the following filters:</i>			
Bonds issued during 2005–2024 period and not by US territories	321,490	2,770,902	
0<offering yield<50% and 50<offering price<150		2,658,486	
0<coupon<20% and offering size>0		2,613,465	
Debt-type="BND" and face value = 100		2,383,107	
Fixed rate coupon		2,382,016	
Purpose of issuance either NEW or REF		2,369,766	
Remove bonds with offering date after maturity date		2,369,657	
Underwriter information available	195,924	2,358,354	
<i>Merged with MSRB trade data:</i>			
Trades within 14 days of offering date			19,168,537
Remove trades of less than \$5,000			19,166,629
Remove trades at prices outside the 50–150 range			19,165,800
Customer buy trades			11,551,952
Sample for new-issue markup analysis	190,225	2,114,399	
<i>Merged with SDC and Bloomberg underwriting spread data:</i>			
Sample with underwriting spread data			138,540

Table 2: Summary statistics of underwriting activity, markups, and spreads

This table reports summary statistics of underwriting activity, markups, and spreads by underwriter type. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. Panel A reports the summary statistics of underwriting activity in an average year. Panel B presents summary statistics on bond-level markups in the new issues market. Markup is defined as the difference between the customer purchase price and the offering price of the bond, expressed in bps of the latter. New issue market trades are defined as those within 14 days of the offering date or when the bond size is fully absorbed by customer buy trades, whichever comes first. The markup for a given bond is the trade-size-weighted average of markups across all new issue trades in the bond. Panel C shows the summary statistics of underwriting spreads (expressed in bps of the total issuance amount). Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum, and Bloomberg. Sample period: 2005–2024.

	Single-state	Small regional	Large regional	National
Panel A: Average annual underwriting activity				
Number of states served	1.00	6.80	27.58	44.19
Number of deals per underwriter	6.82	52.44	252.63	656.22
Volume (\$m) per underwriter	64	619	5,581	19,053
Number of underwriters	113	95	17	8
Market share (% deal count)	4.90	31.83	28.00	35.27
Market share (% volume)	2.24	18.19	29.97	49.60
Panel B: Markups in the new issues market				
Mean markup (bps)	22.32	14.76	13.74	13.30
Median markup (bps)	0.00	0.00	0.00	0.00
StdDev markup (bps)	70.09	61.47	61.54	57.36
Average bond size (\$m)	0.68	0.96	3.32	4.52
Number of bonds	101,506	614,865	553,535	844,493
Panel C: Underwriting spreads				
Mean spreads (bps)	161.23	113.35	90.99	75.18
Median spreads (bps)	129.00	90.00	70.00	60.80
StdDev spreads (bps)	129.58	89.42	91.13	62.25
Average issue size (\$m)	7.08	10.17	38.30	57.62
Number of issues	11,565	65,059	51,292	68,839

Table 3: Time-Varying Decomposition of Underwriter Compensation

The table reports the annual decomposition of underwriter compensation into the retail-distribution rent and the underwriting spread. Effective uninformed volume m_t^U , the conditional markup wedge Δ_t^U , and the rent $\mathcal{R}_t$ are defined in Equations (5)–(7). The underwriting spread S_t is aggregate underwriting spreads per dollar of aggregate issuance. Each row uses only trades issued in the second year of the corresponding two-year estimation window, so annual totals are non-overlapping. The last row is the simple average of the annual values. The sample comprises 8,465,941 customer buy trades in the new issue market for investment-grade rated issues with available underwriting-spread data (109,373 issue-years). Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.

Year	Effective volume m_t^U (%)	Conditional markup wedge		Distribution rent $\mathcal{R}_t = m_t^U \Delta_t^U$ (bps)	Underwriting spread S_t (bps)	Total compensation $C_t = \mathcal{R}_t + S_t$ (bps)	Rent share $\mathcal{R}_t/C_t$ (%)
		Δ_t^U (bps)					
2006	19.87	80.47	15.99	55.80	71.80	22.27	
2007	19.74	85.35	16.85	52.49	69.34	24.30	
2008	19.60	118.87	23.30	56.45	79.75	29.21	
2009	17.37	117.12	20.35	65.14	85.49	23.80	
2010	16.00	102.10	16.33	60.74	77.07	21.19	
2011	15.73	111.08	17.48	57.60	75.07	23.28	
2012	12.92	110.28	14.25	56.12	70.37	20.25	
2013	12.90	111.27	14.35	50.52	64.87	22.12	
2014	12.32	92.38	11.38	48.46	59.84	19.02	
2015	13.39	103.82	13.90	46.41	60.31	23.05	
2016	13.63	100.84	13.74	45.97	59.71	23.01	
2017	12.92	96.49	12.47	43.70	56.17	22.20	
2018	12.77	96.95	12.38	45.21	57.59	21.50	
2019	14.16	81.13	11.49	43.27	54.76	20.98	
2020	9.20	88.00	8.09	44.76	52.85	15.32	
2021	9.63	75.10	7.23	43.92	51.15	14.14	
2022	8.96	63.90	5.73	45.87	51.60	11.10	
2023	9.19	60.76	5.58	46.26	51.84	10.77	
2024	7.36	50.45	3.71	41.83	45.54	8.15	
Average	13.56	91.91	12.87	50.03	62.90	19.77	

Table 4: Sources of the Decline in the Retail-Distribution Rent

The table allocates the change in the retail-distribution rent between the pooled first four years (2006–2009) and last four years (2021–2024) of Table 3 to its two margins, effective uninformed volume m^U and the conditional markup wedge Δ^U , using the identity $\mathcal{R} = m^U \Delta^U$ of Equation (7) and splitting their interaction symmetrically between them (two-factor Shapley allocation; Shorrocks, 2013). The two contributions sum to the total change in the rent. Sample and data sources as in Table 3.

	Pooled level		Contribution to the decline	
	2006–2009	2021–2024	bps	Share (%)
Effective volume m^U (%)	19.05	8.70	−8.46	61.9
Conditional markup wedge Δ^U (bps)	100.47	62.96	−5.21	38.1
Distribution rent $\mathcal{R} = m^U \Delta^U$ (bps)	19.14	5.48	−13.66	100.0

Table 5: Markups and Price Dispersions Around Markup Disclosure Rule

This table reports results for regressions of markups (Panel A) and price dispersion across trades in the same bond (Panel B) in the new issue market around the implementation of the markup disclosure rule. *National* is a dummy variable for the underwriter's geographic breadth, indicating that the underwriter of the issue underwrote issues in 40 or more states in a given year. *Post-Markup Rule* indicates the period after SEC amended rules G-15 and G-30 and FINRA Rule 2232 on markup disclosure become effective May 2018. Bond issue data is from the Mergent municipal fixed income database, trade data is from the MSRB, underwriter spreads (as a percent of issue size) are from SDC Platinum and Bloomberg. For markups, we consider only trades in the new issue market, defined as those within 14 days of the offering date or when the bond size is fully absorbed by customer buy trades, whichever comes first. The sample period is from 2005 through December 2024. Regressions include state and year fixed effects, use of proceeds dummies, and ratings dummies for narrow rating categories. Standard errors are clustered by issue.

Panel A: New Issue Markup %				
	All Bonds	GO Bonds	Issue \$100m+	GO & \$100m+
National	0.010*** (4.03)	0.021*** (6.44)	0.021*** (3.13)	0.034** (2.35)
National x Post Markup Rule	-0.037*** (-9.77)	-0.049*** (-9.41)	-0.029*** (-3.15)	-0.041** (-2.27)
Post Markup Rule	0.005 (1.03)	0.002 (0.25)	0.009 (0.80)	0.033* (1.71)
Bond Controls	Yes	Yes	Yes	Yes
Proceeds Dummies	Yes	Yes	Yes	Yes
Ratings Dummies	Yes	Yes	Yes	Yes
State and Year FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.131	0.119	0.077	0.073
Observations	1,829,323	1,074,288	213,920	65,194
Panel B: Price Dispersion				
	All Bonds	GO Bonds	Issue \$100m+	GO & \$100m+
National	0.019*** (12.01)	0.017*** (8.17)	0.027*** (5.36)	0.025** (2.37)
National x Post Markup Rule	-0.033*** (-11.22)	-0.027*** (-6.85)	-0.030*** (-3.42)	-0.021 (-1.28)
Post Markup Rule	-0.015*** (-3.07) (-0.84)	-0.013** (-2.06) (-0.20)	-0.042*** (-3.20) (-1.57)	-0.066*** (-2.90) (-0.36)
Bond Controls	Yes	Yes	Yes	Yes
Proceeds Dummies	Yes	Yes	Yes	Yes
Ratings Dummies	Yes	Yes	Yes	Yes
State and Year FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.206	0.195	0.245	0.224
Observations	1,082,208	578,586	167,956	48,399

Table 6: Structural Decomposition Around the May 2018 Markup-Disclosure Requirements

This table reports counterfactual quantities from the two-class finite-mixture model in Equations (8)–(10), estimated on the July 2017–April 2018 (pre) and June 2018–March 2019 (post) periods around the May 2018 markup-disclosure requirements (May 2018 is omitted). The specification adds Post Disclosure, Post Disclosure $\times$ National Underwriter, and a linear calendar-month trend to the latent-class equation and both class-specific markup equations. All quantities are means over one common empirical distribution of bond, issue, and transaction characteristics, computed with class priors rather than posteriors. $P(U)$ is the average prior probability of uninformed trading. Effective uninformed volume is prior-weighted customer-buy par divided by represented issue par. The conditional wedge is the prior- and par-weighted positive difference between the fitted uninformed and informed markups. The distribution rent is their product, in basis points of issue par. Panel A retains the observed underwriter mix. Panel B sets National Underwriter to zero or one for every observation, and the National differential is the difference between the two post-minus-pre changes. Delta-method standard errors from the bond-clustered covariance matrix of the joint maximum-likelihood estimate are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels for changes and differences.

	$P(U)$ (percentage points)	Effective uninformed volume (% issue par)	Conditional markup wedge (bps)	Distribution rent (bps issue par)
Panel A: Standardized structural levels				
Pre-disclosure	41.133 (0.412)	14.836 (0.374)	102.555 (2.183)	15.215 (0.478)
Post-disclosure	38.303 (0.438)	12.861 (0.311)	106.641 (2.329)	13.715 (0.463)
Post minus pre	-2.830*** (0.766)	-1.975*** (0.536)	4.086 (3.238)	-1.500** (0.754)
Panel B: Heterogeneity by underwriter reach				
Other underwriters: post minus pre	-2.786*** (0.863)	-2.045*** (0.635)	3.261 (3.343)	-1.715* (0.893)
National underwriters: post minus pre	-2.862*** (0.831)	-1.935*** (0.563)	4.622 (3.690)	-1.381* (0.790)
National differential	-0.077 (0.722)	0.110 (0.512)	1.361 (2.965)	0.333 (0.711)
Post terms in class equation	Yes	Yes	Yes	Yes
Post terms in both markup equations	Yes	Yes	Yes	Yes
Calendar-month trend in both components	Yes	Yes	Yes	Yes
Bond-clustered covariance	Yes	Yes	Yes	Yes
Observations		708,150		
Distinct CUSIPs		117,319		
Distinct issues		9,421		
Log likelihood		10,197,221.0		

Table 7: Local Financial Sophistication and Components of Distribution Rents

This table reports estimates of Equation (15) using state-year data from 2007–2023. The dependent variables are the state-year counterparts of the four quantities in Table ??, computed over bonds issued in each state-year with posterior class probabilities. Direct sophistication is the average share of tax returns reporting taxable interest, dividends, and capital gains, aggregated with fixed national AGI-class weights and expressed in percentage points. Delegated sophistication is the number of SEC- and state-registered investment advisers active in the state per 100,000 residents, scaled so that the coefficient corresponds to ten additional advisers. Both measures are lagged one year and enter jointly. Panel A reports coefficients. Panel B multiplies each coefficient by the interquartile range of within-state deviations in the corresponding measure. All regressions include state and year fixed effects, control for log population, the lagged upper-income return share, and log aggregate issue par, require at least five issues per state-year, and weight by issue count capped at the annual 95th percentile. Standard errors clustered by state are in parentheses and do not account for sampling error in the first-stage mixture estimates. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	Probability uninformed (pp)	Effective uninformed volume (% issue par)	Conditional markup wedge (bps)	Green rent (bps issue par)
Panel A: State-year fixed-effect coefficients				
Direct Sophistication	0.047 (0.665)	0.520 (0.421)	-5.371*** (1.465)	-0.398 (0.448)
Delegated Sophistication	-1.599 (1.675)	0.072 (0.641)	1.308 (2.155)	0.325 (0.716)
Panel B: Effect of a within-state interquartile increase				
Direct Sophistication	0.250 (3.552)	2.779 (2.250)	-28.708*** (7.830)	-2.125 (2.397)
Delegated Sophistication	-1.316 (1.378)	0.059 (0.528)	1.077 (1.773)	0.267 (0.589)
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
State-year controls	Yes	Yes	Yes	Yes
Precision weights	Yes	Yes	Yes	Yes
Observations		841		
States		50		

Table 8: Statistics on the Distribution of Uninformed Trades

This table reports statistics for the relative distribution of the posterior probabilities of informed and uninformed trades computed from the structural model described in Section 4. The sample is split by trade sizes above and below 100,000, and statistics computed separately for each category and year. The first column of each category is the percentage of trades that are uninformed and the second column reports the biserial correlation coefficient between trade size and the indicator for uninformed trade. The third column reports the total variation distance (TVD) between the informed and uninformed probability distributions and represents the symmetric amount of probability mass needed to transfer in order to equalize the two distributions. Below trade sizes of 100,000, the pdf is computed from a discretization using 20 bins and for trade sizes above the pdf is computed using a kernel density estimator with a grid size of 100 and the Parzen kernel.

Uninformed Trade Distribution						
	Trades < 100k			Trades >= 100k		
	% Uninformed	Correlation	TVD	% Uninformed	Correlation	TVD
2006	67.26	-0.119	0.140	38.24	-0.094	0.182
2007	66.61	-0.113	0.129	36.35	-0.075	0.184
2008	68.93	-0.101	0.116	36.74	-0.072	0.205
2009	65.85	-0.137	0.150	31.25	-0.083	0.180
2010	65.41	-0.171	0.171	26.82	-0.068	0.175
2011	67.46	-0.156	0.156	26.41	-0.073	0.219
2012	64.45	-0.154	0.140	22.77	-0.085	0.220
2013	61.03	-0.113	0.097	22.59	-0.074	0.231
2014	54.50	-0.098	0.082	20.68	-0.071	0.227
2015	65.24	-0.109	0.088	24.78	-0.093	0.251
2016	64.87	-0.159	0.146	23.69	-0.087	0.245
2017	60.96	-0.153	0.144	21.03	-0.070	0.255
2018	50.64	-0.127	0.112	20.10	-0.055	0.216
2019	49.39	-0.091	0.076	21.17	-0.057	0.158
2020	41.02	-0.136	0.124	14.49	-0.052	0.131
2021	36.20	-0.098	0.098	13.35	-0.036	0.124
2022	40.46	-0.054	0.099	13.72	-0.047	0.209
2023	43.82	-0.062	0.097	14.46	-0.049	0.269
2024	42.19	-0.070	0.106	14.03	-0.067	0.262
Full Sample	60.37	-0.124	0.119	22.44	-0.068	0.230

Figure 1: Underwriting activity by underwriter type over time

The figure shows the time series variation in underwriting activity. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. The upper panel shows the number of deals (left axis) and the fraction underwritten by each underwriter type (right axis). The lower panel shows the total underwriting volume (in \$ billion, left axis) and the fraction of volume underwritten by each underwriter type (right axis). Data source: Mergent municipal fixed income database. Sample period: 2005–2024.




Figure 2: Time series variation in markups and spreads by underwriter type

This figure presents the average markups and spreads over time by underwriter type. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. Panel A shows the year-by-year average bond-level markups over offering price in the new issue period (within 14 days of the offering date). Panel B plots the year-by-year average underwriting spreads. Data sources: Mergent, MSRB, SDC Platinum, and Bloomberg. Sample period: 2005–2024.






Figure 3: Underwriting Compensation Over Time

The figure plots the time series variation of underwriting revenues. It is underwriting revenue expressed in percentage of issue size. Underwriting revenue includes dollar underwriting fees charged to the issuer and markups charged to investors in the new issue market (i.e., trades within 14 days of issuance or until the total bond amount is fully sold, whichever earlier). Panel A shows the yearly averages across all underwriter types, whereas Panel B shows similar statistics separately for national underwriters, exiting national underwriters (UBS and Citigroup), and other underwriters. National underwriters are those that do business in 40 or more states in a given year. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum, and Bloomberg. Sample period: 2005–2024.




Panel A: All Underwriters




Panel B: By Underwriter Type

Figure 4: Regression Analysis of Markups Over Time

The figure shows the average markups over time after controlling for determinants of markups using the regression model in Equation (13). Panel A plots the Year fixed effects. Panel B plots the National x Year fixed effects. Markup on a customer buy trade is computed as the difference between the price paid and the offering price of the bond, expressed as a percentage of the offering price. The average markup in the new issue market for a given bond is the trade-size-weighted average markups on all trades in that bond within 14 days of issuance or until the total bond amount is fully sold, whichever earlier. Underwriters are classified based on how many states they underwrite in a given year. National underwriters are those that do business in 40 or more states in a given year. Data sources: Mergent municipal fixed income database and MSRB trade data. Sample period: 2005–2024.




Panel A: Year Fixed Effects




Panel B: National x Year Fixed Effects

Figure 5: Regression Analysis of Underwriting Spreads Over Time

The figure shows the average underwriting spreads over time after controlling for determinants of spreads using the regression model in Equation (14). Panel A plots the Year fixed effects. Panel B plots the National x Year fixed effects. Spreads are dollar amounts per \$100 of issuance. Underwriters are classified based on how many states they underwrite in a given year. National underwriters are those that do business in 40 or more states in a given year. Data sources: Mergent municipal fixed income database, SDC Platinum, and Bloomberg. Sample period: 2005–2024.



Panel A: Year Fixed Effects



Panel B: National x Year Fixed Effects

Figure 6: Rolling-Window Probability of the Uninformed Pricing Regime

The figure plots the mean posterior probability that a customer buy trade in the new issue market belongs to the uninformed pricing regime (Equation (2)), averaged over all trades in each overlapping two-year estimation window and dated by the second year of the window. Vertical bars are 95% confidence intervals based on postestimation standard errors clustered by bond. Estimates are based on all trades within 14 days of the offering date or until the bond size is fully absorbed, whichever earlier. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.




Figure 7: Rolling-Window Marginal Mean Markup in the Uninformed Pricing Regime

The figure plots the fitted value of the uninformed-regime outcome equation (Equation (8)), averaged over all trades in each overlapping two-year estimation window and dated by the second year of the window, expressed in bps. Vertical bars are 95% confidence intervals based on postestimation standard errors clustered by bond. The series measures the average level of the latent uninformed pricing schedule; it does not net out the informed-regime counterfactual and does not weight by trade size or posterior probability. Estimates are based on all trades within 14 days of the offering date or until the bond size is fully absorbed, whichever earlier. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.




Figure 8: Underwriting Spreads and the Uninformed Pricing Regime

The figure plots the coefficient on the issue-level underwriting spread in the class-membership equation for each overlapping two-year estimation window, oriented as the log odds of the uninformed pricing regime and dated by the second year of the window. Spreads are measured in percentage points of issuance. Negative estimates indicate that issues with larger issuer-paid spreads have a lower prior probability of being distributed through the uninformed pricing regime, conditional on the model covariates. The coefficient is an equilibrium sorting relation, not a causal marginal effect. Vertical bars are 95% confidence intervals based on standard errors clustered by bond. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.




Figure 9: Distribution Rent and Underwriting Spread by Underwriter Type

The figure plots the market-wide distribution rent (top panel) and the issuance-weighted underwriting spread (bottom panel), in bps of aggregate issuance, computed separately for issues led by single-state, small regional, large regional, and national underwriters (the geographic-reach categories of Section 3). All series use the observation-level posteriors and fitted markup schedules from the same rolling-window models as Table 3. Pooled comparisons are in Online Appendix Table A.5. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.





Figure 10: Relative Distributions of Uninformed/Informed Trades

The figure plots two versions of the distributional overlap between uninformed and informed trades below 100,000. Uninformed and informed trades are defined as the posterior probabilities generated from the structural model in Section 4. Panel A plots histograms for each distribution using bins of width 5,000. Panel B implements the point-by-point multiple testing of the CDFs as in [Goldman and Kaplan \(2018\)](#), with regions where equality can be rejected highlighted by the bold black line on the x-axis. Sample period: 2005–2024.



A Online Appendix

Table A.1: Identification and Posterior Separation in the Finite-Mixture Model

This table reports identification and classification diagnostics for each overlapping two-year finite-mixture estimation window. Each row labels the two-year window by its second year. The conditional residual standard deviations, $\hat{\sigma}_I$ and $\hat{\sigma}_U$, are posterior-weighted and expressed in basis points. Following GHS (2007), the lower-dispersion component is labeled informed, so the identifying ordering is $\hat{\sigma}_I < \hat{\sigma}_U$. Because the raw dispersion ratio becomes extremely large when the informed component approaches a mass point at the reoffering price, the table reports $\log_{10}(\hat{\sigma}_U/\hat{\sigma}_I)$; the underlying ratio is no smaller than 2,652 in any window. Normalized posterior entropy equals the mean binary entropy divided by $\log 2$ and multiplied by 100, so zero denotes deterministic classification and 100 denotes maximal ambiguity. A posterior is decisive when $\hat{P}(U) \leq 0.05$ or $\hat{P}(U) \geq 0.95$, and ambiguous when $0.25 < \hat{P}(U) < 0.75$. Beginning in 2017, the informed component is estimated close to a mass point at the reoffering price, consistent with the growing concentration of informed trades at zero markup. All diagnostics are calculated on the exact full two-year FMM estimation sample, rather than the non-overlapping endpoint-year accounting sample. No structural model is re-estimated for this table.

Endpoint year	Estimation N (000s)	Conditional-dispersion separation			Posterior separation		
		$\hat{\sigma}_I$ (bps)	$\hat{\sigma}_U$ (bps)	$\log_{10}(\hat{\sigma}_U/\hat{\sigma}_I)$	Normalized entropy (% max)	Decisive posteriors (%)	Ambiguous posteriors (%)
2006	684	0.032	85.67	3.42	0.489	99.619	0.145
2007	675	0.030	85.41	3.45	0.333	99.803	0.074
2008	881	0.031	103.42	3.53	0.335	99.745	0.074
2009	1269	0.032	106.32	3.52	0.316	99.734	0.068
2010	1241	0.032	97.89	3.49	0.267	99.813	0.049
2011	976	0.030	89.97	3.47	0.214	99.878	0.036
2012	965	0.030	91.14	3.48	0.181	99.915	0.026
2013	933	0.029	93.21	3.50	0.186	99.907	0.021
2014	866	0.029	88.28	3.48	0.202	99.889	0.025
2015	1008	0.028	83.51	3.48	0.223	99.872	0.033
2016	1148	0.018	86.69	3.67	0.859	99.055	0.301
2017	1127	< 0.001	88.46	10.70	0.007	99.992	0.005
2018	930	< 0.001	86.07	10.59	0.002	99.997	0.001
2019	762	< 0.001	80.64	7.96	0.001	99.999	< 0.001
2020	743	< 0.001	79.19	6.40	0.003	99.999	< 0.001
2021	678	< 0.001	77.14	11.80	< 0.001	100.000	< 0.001
2022	619	< 0.001	93.13	6.96	0.008	99.984	< 0.001
2023	646	< 0.001	107.11	7.74	0.001	100.000	< 0.001
2024	740	< 0.001	95.63	6.77	0.006	99.994	< 0.001
Extrema across windows					0.859	99.055	0.301

Table A.2: Distribution of Days Since Offering Until Fully Sold

This table reports the distribution of the number of days between the offering date and the date the bond is fully sold. The sample period is from 2005 to 2024.

Number of days	% of bonds	Cumulative % of bonds
<0	0.53	0.53
0	18.06	18.60
1	34.37	52.97
2	3.96	56.92
3	2.00	58.92
4	1.95	60.87
5	2.07	62.93
6	2.39	65.32
7	2.31	67.63
8-14	7.96	75.60
15-30	9.90	85.50
31-60	6.36	91.86
61-90	2.00	93.86
91-120	0.80	94.66
121-180	0.61	95.27
180+	4.73	100.00

Table A.3: Classification Regression Coefficients

This table reports coefficient estimates and t -statistics (in parentheses) for the latent-variable logit classification equation of the structural model in Section 2.2, for the two-year estimation windows ending in 2008, 2012, 2017, 2020, and 2024. The full system is estimated by maximum likelihood, with starting values from an expectation-maximization step. t -statistics are based on standard errors clustered by bond. *, **, and *** denote significance at the 5%, 1%, and 0.1% levels.

	2008	2012	2017	2020	2024
Ln(Trade Size)	-0.359*** (-45.52)	-0.372*** (-48.74)	-0.301*** (-32.14)	-0.203*** (-25.44)	-0.241*** (-40.92)
Ln(Issue Amount)	-0.184*** (-11.32)	-0.130*** (-8.60)	0.060 (0.51)	0.001 (0.07)	-0.086*** (-6.15)
Days Since Offer	0.427*** (33.11)	0.416*** (50.04)	0.556*** (43.46)	0.538*** (53.05)	0.649*** (83.91)
Fraction of Bond Sold	0.820*** (17.00)	0.702*** (18.21)	0.577*** (7.02)	0.448*** (13.76)	1.270*** (36.58)
Issued at Premium	-0.399*** (-11.92)	-0.666*** (-18.32)	-0.721*** (-3.76)	-0.552*** (-16.71)	-1.033*** (-30.31)
Bond Size	0.000 (1.69)	0.000* (2.21)	-0.000 (-0.00)	-0.000 (-0.64)	0.000* (2.27)
Coupon	-0.038 (-1.02)	0.039 (1.79)	-0.112 (-1.55)	0.063*** (4.54)	0.019 (0.79)
Maturity	0.101*** (10.44)	0.144*** (12.52)	0.106*** (13.02)	0.073*** (16.03)	0.019* (2.10)
Maturity Sqr	-0.001*** (-4.75)	-0.002*** (-6.77)	-0.001 (-1.48)	-0.000*** (-5.55)	0.001*** (3.48)
Callable Y/N	0.385*** (6.93)	0.210*** (4.50)	0.273*** (3.55)	0.118** (3.08)	0.088* (2.03)
National Und.	0.053 (1.27)	-0.029 (-0.73)	-0.149*** (-4.08)	-0.355*** (-13.52)	-0.041 (-1.35)
Negotiated	-0.451*** (-9.88)	-0.153*** (-3.98)	-0.286*** (-4.34)	-0.599*** (-18.85)	-1.115*** (-28.13)
Bond type: Other (vs. GO)	0.040 (0.86)	-0.310*** (-6.59)	-0.320*** (-6.01)	-0.481*** (-12.95)	-0.440*** (-9.25)
Bond type: Revenue (vs. GO)	0.041 (0.73)	-0.415*** (-7.58)	-0.260** (-2.74)	-0.248*** (-5.03)	-0.281*** (-5.33)
Insured (Y/N)	0.385*** (9.57)	0.745*** (13.17)	0.400*** (7.35)	0.034 (0.77)	-0.024 (-0.40)
Credit Enhancement	0.375*** (8.61)	-0.067 (-1.29)	-0.019 (-0.29)	0.099* (2.37)	0.093 (1.84)
Bank Qualified	-0.492*** (-10.60)	-0.088* (-1.98)	0.164 (1.58)	0.281*** (8.32)	-0.019 (-0.39)
First-time Issuer	0.169*** (3.57)	0.247*** (4.63)	-0.172 (-1.66)	0.160*** (3.54)	0.021 (0.36)
Issue Frequency	-0.005*** (-3.54)	-0.002*** (-5.78)	-0.002*** (-8.55)	-0.000 (-0.96)	-0.000*** (-3.45)
Underwriter Spread	-1.401*** (-16.22)	-0.858*** (-12.67)	-0.541*** (-10.06)	-0.540*** (-12.10)	-0.445*** (-13.29)
Ratings Dummies	Yes	Yes	Yes	Yes	Yes
Proceeds Dummies	Yes	Yes	Yes	Yes	Yes
N	881,294	965,032	1,127,232	742,727	739,842

Table A.4: Markup Regression Coefficients for Uninformed Trades

This table reports coefficient estimates and t -statistics (in parentheses) for the uninformed-regime markup equation (Equation (8)) of the structural model in Section 2.2, for the two-year estimation windows ending in 2008, 2012, 2017, 2020, and 2024. The full system is estimated by maximum likelihood, with starting values from an expectation-maximization step. t -statistics are based on standard errors clustered by bond. *, **, and *** denote significance at the 5%, 1%, and 0.1% levels.

	2008	2012	2017	2020	2024
Ln(Trade Size)	2221.435*** (20.79)	1503.686*** (20.08)	1325.688*** (11.88)	1481.401*** (16.04)	-526.757*** (-5.23)
Ln(Issue Amount)	0.037* (2.48)	-0.019* (-2.08)	0.010 (0.19)	-0.006 (-0.71)	-0.035*** (-4.26)
Days Since Offer	-0.005 (-1.74)	0.012*** (6.77)	-0.001 (-0.32)	0.028*** (19.89)	-0.003 (-1.61)
Fraction of Bond Sold	0.035 (0.72)	0.137*** (6.22)	0.181*** (3.55)	-0.080*** (-4.49)	-0.058** (-2.64)
Issued at Premium	-0.829*** (-20.96)	-0.563*** (-25.12)	-0.339*** (-4.64)	-0.374*** (-21.04)	-0.569*** (-25.61)
Bond Size	-0.000 (-0.42)	-0.000 (-1.01)	-0.000 (-0.11)	-0.000** (-2.77)	-0.000 (-0.65)
Coupon	0.181*** (5.99)	0.044*** (3.63)	-0.048 (-1.52)	-0.066*** (-9.93)	0.009 (0.62)
Maturity	0.077*** (7.36)	0.116*** (12.98)	0.103*** (17.09)	0.083*** (25.67)	0.037*** (6.56)
Maturity Sqr	-0.001*** (-3.41)	-0.002*** (-7.50)	-0.002*** (-5.99)	-0.001*** (-9.29)	0.000 (0.17)
Callable Y/N	-0.158** (-2.58)	0.021 (0.59)	0.271*** (9.96)	0.017 (0.90)	0.011 (0.38)
National Und.	-0.007 (-0.20)	-0.117*** (-6.00)	-0.124*** (-7.55)	-0.033* (-2.16)	0.048* (2.45)
Negotiated	0.385*** (11.86)	0.251*** (13.04)	0.319*** (13.13)	0.480*** (30.61)	0.605*** (27.15)
Bond type: Other (vs. GO)	0.054 (1.68)	0.029 (1.11)	0.082*** (3.32)	0.072*** (3.37)	-0.036 (-1.46)
Bond type: Revenue (vs. GO)	0.074* (2.01)	0.060 (1.96)	0.036 (1.29)	0.118*** (4.60)	0.017 (0.58)
Insured (Y/N)	0.115** (3.16)	0.153*** (6.69)	0.092** (3.11)	-0.011 (-0.38)	-0.061 (-1.84)
Credit Enhancement	-0.029 (-0.83)	-0.008 (-0.25)	0.026 (0.97)	-0.008 (-0.39)	0.050 (1.38)
Bank Qualified	-0.170*** (-4.55)	-0.294*** (-12.03)	-0.140** (-3.20)	-0.017 (-1.11)	-0.019 (-0.90)
Constant	-1.025*** (-3.53)	0.426** (2.69)	-0.102 (-0.12)	0.161 (1.20)	0.917*** (5.97)
Ratings Dummies	Yes	Yes	Yes	Yes	Yes
Proceeds Dummies	Yes	Yes	Yes	Yes	Yes
N	881,294	965,032	1,127,232	742,727	739,842

Table A.5: Decomposition of Underwriter Compensation by Underwriter Type

This table reports the decomposition of Table 3 computed separately within each geographic-reach category of Section 3, pooling the first four years of the decomposition (2006–2009) and the last four (2021–2024). $P(U)$ is the mean posterior probability of the uninformed pricing regime; the remaining objects follow Section 2. All values use the same fitted structural models as the main analysis. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.

	$P(U)$ (%)		m^U (%)		Δ^U (bps)		$\mathcal{R}$ (bps)		S (bps)		C (bps)	
	06–09	21–24	06–09	21–24	06–09	21–24	06–09	21–24	06–09	21–24	06–09	21–24
Underwriter type												
Single-state	63.28	67.23	21.30	26.07	124.78	100.06	26.58	26.09	66.23	111.70	92.80	137.79
Small regional	61.71	34.90	24.98	15.42	106.42	69.89	26.58	10.78	69.04	71.43	95.62	82.21
Large regional	54.96	26.30	19.82	9.73	100.40	63.82	19.90	6.21	58.80	45.61	78.69	51.82
National	55.09	20.36	16.87	7.65	97.13	60.50	16.38	4.63	53.82	41.23	70.20	45.86

Table A.6: Classification Leverage, Offering Method, and Conservative Retail-Order-Period Bounds

This table evaluates whether the decline in the distribution rent is mechanically generated by the construction of the posterior classifications (Panel A), the mix of competitive and negotiated offerings (Panel B), or early negotiated-sale allocations at the reoffering price (Panel C). All calculations use the observation-level predictions from the same fitted models as Table 3. Early and late statistics pool 2006–2009 and 2021–2024, respectively. In Panel A, the density-only posterior assigns every trade in a window the window-mean class prior while preserving the fitted markup densities (Equation (12)); the spread-neutralized posterior fixes every issue’s underwriting spread at its within-window mean; the hard classification assigns each trade entirely to its more likely regime based on the posterior probability; maximum annual deviations are relative to the baseline posterior. Panel C excludes negotiated issues whose customer buy volume executed through event day 1 (including when-issued trades) within one bp of the reoffering price is at least the stated share of the issue’s observed customer buy volume; these bounds intentionally over-exclude and their levels are not comparable to the market-wide estimates. $P(U)$, Δ^U , and the distribution rent $\mathcal{R}$ follow Section 2, with $\mathcal{R}$ in bps of aggregate issuance. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.

Panel A: Alternative classification constructions

Classification	$P(U)$ (%)			$\mathcal{R}$ (bps)			Maximum annual deviation		
	06–09	21–24	06–09	21–24	Change	$ \Delta P(U) $ (pp)	$ \Delta \mathcal{R} $ (bps)		
Baseline posterior	56.42	24.11	19.14	5.48	–13.66	0.000	0.000		
Density-only posterior	56.43	24.11	19.20	5.47	–13.73	0.027	0.222		
Spread fixed at window mean	56.43	24.11	19.14	5.48	–13.66	0.016	0.048		
Hard classification	56.40	24.11	19.15	5.48	–13.66	0.032	0.073		

Panel B: Offering method

Offering method	$P(U)$ (%)			Δ^U (bps)			$\mathcal{R}$ (bps)		
	06–09	21–24	06–09	21–24	Change	$ \Delta P(U) $ (pp)	$ \Delta \mathcal{R} $ (bps)		
All offerings	56.42	24.11	100.47	62.96	19.14	5.48			
Competitive	65.08	50.31	55.07	35.63	17.12	9.01			
Negotiated	54.54	15.56	118.04	94.65	19.56	4.68			

Panel C: Conservative retail-order-period bounds within negotiated offerings

Negotiated-issue sample	Remaining issues			$P(U)$ (%)			Δ^U (bps)			$\mathcal{R}$ (bps)		
	06–09	21–24	14,118	06–09	21–24	Change	06–09	21–24	Change	06–09	21–24	Change
All negotiated issues	12,801	14,118	54.54	15.56	118.04	94.65	19.56	4.68				
Exclude early near-zero share $\geq 10\%$	1,535	368	73.89	35.39	116.98	115.48	42.98	18.67				
Exclude early near-zero share $\geq 25\%$	2,285	553	77.72	50.66	115.17	114.94	49.50	22.64				
Exclude early near-zero share $\geq 50\%$	3,921	1,122	76.66	58.41	114.44	113.43	50.60	30.24				

Table A.7: Structural Decomposition Across New-Issue Trade Windows

This table recomputes the structural objects in Table 3 over alternative trade horizons, using the same observation-level posterior probabilities and fitted informed and uninformed markup schedules from the baseline rolling-window estimation. Early and late statistics pool 2006–2009 and 2021–2024, respectively, by summing the underlying numerators and denominators. Calendar event day is measured relative to the offering date. The first-five-trading-dates row ranks distinct observed transaction dates within each CUSIP, including when-issued dates when present in the baseline sample, and is the closest available counterpart to the first-five-trading-day horizon in GHS (2007). For every narrow window, effective uninformed volume and the distribution rent retain the full aggregate issuance denominator of the corresponding years; represented issuance and customer-buy coverage are reported separately so that later windows do not acquire a mechanically smaller denominator. $P(U)$ is the transaction-weighted mean posterior probability of the uninformed regime. Effective uninformed volume is posterior-weighted customer-buy volume divided by aggregate issuance. The conditional wedge is the posterior- and trade-size-weighted difference between the fitted uninformed and informed markup schedules. The distribution rent is their product, expressed in bps of aggregate issuance. Panel B reports the GHS-style hard-classified markup differential per uninformed transaction, the observed trade-size-weighted raw markup, and customer-buy volume relative to aggregate issuance.

Panel A: Rent decomposition

Trade horizon	$P(U)$		Effective uninformed volume		Conditional wedge		Distribution rent	
	2006–09	2021–24	2021–24 (% of issuance)		2021–24 (bps)		2021–24 (bps of issuance)	
Full baseline: absorbed or 14 days	0.564	0.241	19.05	8.70	100.47	62.96	19.14	5.48
Through event day 1 (incl. when-issued)	0.283	0.067	9.22	3.35	94.40	68.82	8.70	2.30
Calendar event days 0-1	0.283	0.066	8.99	3.34	93.95	68.83	8.44	2.30
Calendar event days 2-5	0.811	0.756	3.88	2.10	112.16	70.03	4.35	1.47
Calendar event days 2-14	0.880	0.839	9.83	5.36	106.16	59.30	10.44	3.18
Calendar event days 1-5	0.570	0.184	9.12	4.85	100.45	67.06	9.16	3.25
First five observed trading dates	0.519	0.225	17.58	8.38	96.72	62.48	17.00	5.23

Panel B: Transaction-level and coverage diagnostics

Trade horizon	Hard-classified MLOT/trade		Observed raw markup		Customer-buy coverage	
	2006–09	2021–24 (bps)	2021–24 (bps)	2021–24 (% of issuance)	2021–24 (% of issuance)	2021–24 (% of issuance)
Full baseline: absorbed or 14 days	157.74	102.75	7.64	1.95	80.65	95.01
Through event day 1 (incl. when-issued)	158.84	108.25	-0.41	0.80	67.63	87.94
Calendar event days 0-1	158.88	108.27	-0.52	0.80	66.29	87.87
Calendar event days 2-5	165.95	117.31	46.16	20.89	6.51	3.59
Calendar event days 2-14	157.34	101.26	49.44	16.27	13.02	7.07
Calendar event days 1-5	162.60	113.79	8.88	2.06	34.57	66.20
First five observed trading dates	154.15	100.87	5.79	1.83	79.11	94.68

Table A.8: Issue-Level Compensation Statistics

This table reports unweighted issue-level counterparts to the market-wide decomposition in Table 3. The issue-mean rent averages, across issues in each year, an issue’s posterior-weighted money left on the table divided by its issuance amount; the bond-mean rent is the corresponding average across bonds; the issue-mean spread is the unweighted average underwriting spread; and the issue-mean rent share averages issue-level rent shares of total compensation. Because large issues generate lower rent per dollar and pay lower percentage spreads, these unweighted means are higher than the market-wide aggregates in Table 3. Sample as in Table 3. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.

Year	Issue-mean rent (bps)	Bond-mean rent (bps)	Issue-mean spread (bps)	Issue-mean rent share (%)
2006	21.70	21.61	73.20	19.94
2007	22.30	21.44	73.40	19.82
2008	28.70	24.70	76.93	23.09
2009	27.90	25.69	85.86	20.20
2010	24.38	23.36	82.43	18.78
2011	26.58	24.94	84.43	20.50
2012	22.12	21.13	79.51	18.21
2013	22.41	21.30	77.38	18.71
2014	20.08	18.96	74.66	17.81
2015	23.15	22.42	70.90	20.89
2016	20.54	18.91	69.38	19.25
2017	19.87	18.45	75.40	17.90
2018	18.05	17.04	81.34	16.03
2019	16.93	15.79	75.07	15.90
2020	14.87	13.19	77.97	13.44
2021	11.70	10.54	74.96	11.36
2022	12.60	9.06	84.43	10.20
2023	11.76	9.29	85.19	9.70
2024	8.59	7.41	78.44	8.61

Table A.9: Long-Run Changes in Sophistication and Components of Distribution Rents

This table reports long-difference counterparts of the regressions in Table 7, where all variables are defined. For each state, the variables are averaged over 2007–2010 and over 2020–2023, and the change in each dependent variable is regressed on the changes in one-year-lagged direct and delegated sophistication, which enter jointly, and on the changes in log population, the lagged upper-income return share, and log aggregate issuance. Panel A weights states equally and uses HC3 standard errors. Panel B weights states by issue count, capped at the cross-state 95th percentile. The IQR rows multiply each coefficient by the cross-state interquartile range of the corresponding long-run change. Standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	Probability uninformed (pp)	Effective uninformed volume (% issue par)	Conditional markup wedge (bps)	Green rent (bps issue par)
Panel A: Equal-state long differences				
Change in direct sophistication	-0.850 (1.216)	-0.896 (0.861)	-4.474 (2.880)	-1.661* (0.862)
Change in delegated sophistication	-1.414 (1.523)	-0.151 (0.506)	-1.342 (1.893)	-0.256 (0.535)
Change in direct sophistication – IQR effect	-1.456 (2.082)	-1.534 (1.474)	-7.663 (4.933)	-2.846* (1.476)
Change in delegated sophistication – IQR effect	-1.121 (1.208)	-0.120 (0.401)	-1.064 (1.501)	-0.203 (0.425)
Panel B: Issue-count-weighted long differences				
Change in direct sophistication	0.213 (1.347)	0.229 (0.800)	-6.198** (2.982)	-0.903 (0.845)
Change in delegated sophistication	-2.056 (2.266)	-0.535 (0.774)	0.476 (2.480)	-0.566 (1.023)
Change in direct sophistication – IQR effect	0.365 (2.308)	0.391 (1.370)	-10.616** (5.107)	-1.546 (1.447)
Change in delegated sophistication – IQR effect	-1.630 (1.797)	-0.424 (0.613)	0.378 (1.966)	-0.449 (0.811)
States	50			
Periods compared	2007–2010 versus 2020–2023			

Figure A.1: Composition of Underwriters (By State)

Figure shows the fraction of each underwriter type in an average year for each given state. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. Data source: Mergent municipal fixed income database. Sample period: 2005–2024.




Figure A.2: Fraction of Deals (By Underwriter Type and State)

Figure shows the fraction of deals underwritten by each underwriter type in an average year for each given state. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. Data source: Mergent municipal fixed income database. Sample period: 2005–2024.




Figure A.3: Fraction of Volume (By Underwriter Type and State)

Figure shows the fraction of volume underwritten by each underwriter type in an average year for each given state. Underwriters are classified based on how many states they underwrite in a given year: “single state” if one state, “small regional” if between 2 and 19 states, “large regional” if between 20 and 39, and “national” if 40 and above. Data source: Mergent municipal fixed income database. Sample period: 2005–2024.




Figure A.4: Probability of Trading with Uninformed Investors: Revenue Bonds

The figure plots the mean posterior probability that a customer buy trade in the new issue market belongs to the uninformed pricing regime (Equation (2)), as in Figure 6, with the model estimated on rated revenue bonds only; vertical bars are standard errors. Estimates are based on overlapping two-year windows dated by the second year of the window. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.




Figure A.5: Markup on Uninformed Trades: Revenue Bonds

The figure plots the latent class mean markups and their standard errors from the mixture model with two latent classes for uninformed trades. Markups for informed trades are effectively zero for all periods. Estimates are based on all trades in the new issue market in rated Revenue bonds in rolling two-year periods beginning in 2006. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum, and Bloomberg. Sample period: 2005–2024.




Figure A.6: Effects of National Underwriter on the Probability of Uninformed Trades

The figure plots the effects of the National Underwriter dummy variable on the probability of a trade with uninformed investors. The estimate is from the classification regression of the mixture model with informed and uninformed investors. National underwriters are those that underwrite bonds in 40 or more states in a given year. Estimates are based on all trades within 14 days of the offering date or until the bond size is fully absorbed, whichever earlier. The model is estimated separately for overlapping two-year periods beginning in 2005-2006. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.




Figure A.7: Effects of National Underwriter on the Markups on Uninformed Trades

The figure plots the coefficient of the National Underwriter dummy in the markup outcome equation in the uninformed regime. National underwriters are those that underwrite bonds in 40 or more states in a given year. Markups are expressed in percentage points. Estimates are based on all trades within 14 days of the offering date or until the bond size is fully absorbed, whichever earlier. The model is estimated separately for overlapping two-year periods beginning in 2005-2006. Data sources: Mergent municipal fixed income database, MSRB trade data, SDC Platinum and Bloomberg. Sample period: 2005–2024.

