

Information Leakage in US Credit: Evidence from All-to-All RFQs

Filippo Caretti*

Harveer Mahajan[†]

Ted Husveth[‡]

May 2026

Abstract

In this paper, we estimate the relative market impact of two widely used trading protocols within the U.S. corporate bond market: All-to-All (A2A) request-for-quote (RFQ) trading and Voice (bilateral) processed trades. Specifically, we implement a generalized linear model (GLM) with a balanced weighting scheme to estimate the conditional difference in cumulative spread changes between RFQ acceptance and public reporting on TRACE. Our analysis leverages Tradeweb’s internal RFQ datasets alongside Ai-Price, a proprietary pricing algorithm that provides dense and consistent intraday price estimates. We focus exclusively on investment grade (IG) trades with notional values exceeding \$1 million. The motivation for this analysis stems from Tradeweb’s central role in an increasingly electronic corporate bond market, where rising transparency has potentially amplified the cost of systemic inefficiencies, such as the broadcasting of trading intent. Our findings indicate that the A2A trading protocol, although proven to have better execution quality, is statistically associated with relatively greater information leakage compared to equivalent bilateral trades. This result supports the development of trading protocols such as SNAP+, which aim to preserve competitive pricing while reducing the diffusion of sensitive trading information inherent to one-to-many RFQ structures.

*filippo.caretti@tradeweb.com

†harveer.mahajan@tradeweb.com

‡ted.husveth@tradeweb.com

1 Introduction

Corporate bond markets have traditionally been characterized by an over-the-counter (OTC) market structure and limited transparency. Over time, regulatory initiatives, technological advancements, and new trading protocols have improved the timeliness, granularity and accessibility of transaction data. At the same time, credit market participants have become more sophisticated, though corporate bond markets still lag other asset classes in terms of automation and efficiency. As markets continue to evolve, so too do opportunities that arise from inefficiencies related to how trading intentions are revealed.

This paper focuses on one of the central risks in increasingly electronic corporate bond markets: **information leakage**. We analyze how different trading protocols affect the visibility of clients' trading interests to the broader market, and how that visibility can shape short-term price dynamics.

We focus on two dominant protocols: All-to-All request-for-quotes (A2A RFQs) and bilateral trading or Voice Processing (Voice). A2A RFQs enable clients to broadcast a trading intent to an extensive set of potential counterparties, including anonymous buy-side participants. This approach offers deeper liquidity and more competitive pricing, often resulting in tighter bid-offer spreads. However, this broader dissemination of intent introduces a structural drawback: increased exposure may lead to information leakage, amplifying same-direction short-term price pressure. In contrast, Voice trading involves bilateral negotiations with a single dealer, limiting information spillover, but also narrowing the pool of liquidity providers, which may result in less competitive pricing and higher adverse selection risk.

While prior research highlights the efficiency benefits of A2A RFQs, the central objective of this paper is only to *quantify and compare the degree of information leakage between A2A RFQs and Voice trades*. We argue that as electronification continues, the trade-off between transparency, competition, and information control will become increasingly important for both buy-side and sell-side participants. Our analysis contributes to this debate by providing empirical evidence on how disclosure through A2A protocols shapes short-term price dynamics, relative to more traditional Voice execution.

2 Data

Our analysis relies on Tradeweb's platform data covering client A2A RFQ and Voice trades from January 1, 2025 to April 30, 2026. We restrict the sample to transactions with notional volume greater than \$1 million, as larger trades are more likely to reflect meaningful trading intent and contribute to price discovery.

To ensure accurate comparability across instruments, we limit the universe to **investment-grade (IG)** corporate bonds quoted in spread terms. After implementing these filters, the sample is still large enough to make robust inferences. Although there are relatively equal amounts of data for both A2A and Voice trades, we structure the analysis under a *balanced design* to account for any differences in class representation

during the observed period and to mitigate potential biases in estimated coefficients arising from the empirical sample.

For pricing, we rely on Tradeweb **Ai-Price**, which provides dense intraday mid-price levels for corporate bonds. Ai-Price allows us to align realized trade data with consistent price estimates and ensures minimal loss of observations due to missing data, thereby preserving the integrity of the event study framework.

3 Methodology

Access to Tradeweb’s RFQ data provides a unique opportunity to study the dynamics of information leakage in U.S. corporate bond trading. Specifically, we can isolate the interval between RFQ acceptance and TRACE publication. During this window, price movements are not yet observable to the broader market and therefore provide a natural setting to measure information leakage. This framework lends itself to a treatment-control design, where we compare the cumulative spread dynamics of A2A versus Voice trades.

3.1 Outcome Variable

Our outcome variable is the **cumulative spread change** measured over successive five-minute intervals between RFQ acceptance and TRACE publication. By adopting a relatively high-frequency evaluation window, we can construct a *price impact curve* that traces how spreads evolve dynamically after an RFQ is executed.

Formally, let $\{M_{t_k}\}_{k=0}^K$ denote the sequence of mid prices (from Ai-Price - in spread terms) aligned into five-minute buckets, starting at trade time t_0 . The incremental log return is defined as:

$$\Delta r_{t_k} = \log\left(\frac{M_{t_k}}{M_{t_{k-1}}}\right),$$

and the cumulative spread change up to horizon τ is given by:

$$r_{t_0 \rightarrow t_0 + \tau} = \sum_{k=1}^{\tau} \Delta r_{t_k}.$$

This construction yields a trajectory of spread responses to trading events, indexed by elapsed time since the RFQ. Trades without a valid Ai-Price series over the relevant horizon are excluded. By design, this framework isolates the effect of information leakage prior to public TRACE disclosure, since the trade is not yet observable to the wider market.

3.2 Hypothesis

Our study is guided by the following hypothesis:

Hypothesis (H_1): A2A RFQs generate greater same-direction price impact relative to Voice trades during the window between trade execution and TRACE publication, due to the wider dissemination of trading intent.

The corresponding null hypothesis is:

Null Hypothesis (H_0): There is no systematic difference in cumulative spread changes between A2A and Voice trades during the pre-TRACE window.

Testing this hypothesis amounts to estimating the conditional difference in cumulative spread trajectories across the two trade types. A statistically significant positive coefficient on the A2A indicator would provide evidence in favor of H_1 , consistent with the notion that broadcasting trading intent accelerates same-direction market movements.

3.3 Balanced Framework

Having defined our outcome variable and the hypothesis of our study, we turn our attention to the estimation methodology. We propose a treatment-control framework where we consider the following A2A dummy indicator as a factor associated with our outcome variable:

$$\mathbb{1}_{A2A} = \begin{cases} 1 & \text{if } x \in \text{A2A} \\ 0 & \text{if } x \in \text{Voice} \end{cases}$$

To account for trade direction, we estimate the model separately for buyer- and seller-initiated transactions. This distinction is crucial, as the price impact dynamics associated with purchases and sales are inherently asymmetric. Buyer-initiated trades typically exert upward price pressure (spread tightening), while seller-initiated trades generate downward price pressure (spread widening). Pooling buys and sells would therefore conflate these opposing effects, obscuring the true relationship between execution protocol and transaction-level price impact. By conditioning our analysis on trade direction, we isolate within-direction variation attributable to the execution protocol.

Control Variables: Trade Volume and Liquidity Score

To ensure that the estimated effect of the execution protocol differences (A2A vs. Voice) is not confounded by broader market conditions or trade-specific characteristics, we introduce explicit controls for trade volume and liquidity - two well-established determinants of transaction cost and price impact.

We leverage Tradeweb’s proprietary Liquidity Score as a proxy for market activity and trading intensity associated with a given bond. This score is mostly based on 90-day rolling averages of TRACE-reported

trading volumes, which are subsequently ranked into deciles to obtain a standardized and comparable measure across instruments and time. This deciled transformation mitigates the influence of extreme observations and accommodates the heavy-tailed nature of the volume distribution in corporate bond markets. Conceptually, this measure captures both cross-sectional and temporal variation in market liquidity, thereby allowing us to control for systematic differences in trading conditions that could otherwise bias the estimated treatment effect.

We further control for trade size, as empirical evidence consistently documents a positive relationship between transaction size and execution costs. Larger trades tend to generate greater price impact because of limited market depth and increased inventory risk. Moreover, empirical observations suggest a degree of non-linearity; accordingly, we include the natural logarithm of accepted trade size as a covariate, allowing for diminishing marginal effects of trade volume on the outcome variable.

Modeling the Objective

Our modeling objective is to estimate the conditional distribution of the outcome variable - typically a short-horizon price response or transaction cost - given both the execution protocol and the relevant covariates. This framework allows us to isolate the incremental effect of A2A execution while partialling out observable determinants of liquidity and trade size. In other words, we seek to isolate the component of price impact variation that is orthogonal to systematic liquidity and volume effects.

Formally, we specify a generalized linear model (GLM) in which the conditional distribution of the outcome variable is Gaussian with an identity link. For a trade observed at time t_0 , with an outcome variable measured over horizon τ , the model is given by:

$$r_{t_0 \rightarrow t_0 + \tau} \sim \mathcal{N} \left(\beta_{0,\tau} + \beta_{A2A,\tau} \mathbb{1}_{A2A} + \beta_{vol,\tau} \log(\text{AcceptedSize}) + \beta_{liq,\tau} \text{LiquidityScore}, \sigma^2 \right)$$

In this formulation, $\beta_{A2A,\tau}$ captures the conditional difference in the expected outcome between A2A and Voice executions, holding liquidity and trade size constant. The parameters $\beta_{vol,\tau}$ and $\beta_{liq,\tau}$ quantify, respectively, the sensitivity of price impact to trade volume and prevailing liquidity conditions.

Weighted GLM for Class Balance

In principle, our empirical sample of IG trades with volume \geq \$1 million does not exhibit major class imbalance between A2A and Voice observations. Nevertheless, to guard against small-sample distortions, we fit a GLM with a *daily class-balancing* weighting scheme.

For each trading day d , let $N_{0,d}$ and $N_{1,d}$ denote the number of Voice and A2A trades, respectively, with

$N_d = N_{0,d} + N_{1,d}$. For any observation i belonging to class $c_{i,d} \in \{0, 1\}$ on day d , we define the observation-specific weight as:

$$w_{i,d} = \begin{cases} \frac{N_d}{2N_{0,d}}, & \text{if } c_{i,d} = 0 \text{ (Voice)} \\ \frac{N_d}{2N_{1,d}}, & \text{if } c_{i,d} = 1 \text{ (A2A)} \end{cases}$$

This ensures that within each trading day, both classes contribute equally to the weighted likelihood:

$$\sum_{i \in d: c=0} w_{i,d} = \sum_{i \in d: c=1} w_{i,d} = \frac{N_d}{2}.$$

Let $r_{i,d}^{(\tau)}$ denote the cumulative spread change for observation i on day d , measured from the RFQ acceptance time t_0 to $t_{0+\tau}$. The weighted log-likelihood of the model can then be written as:

$$\ell(\beta_\tau) = \sum_d \sum_{i \in d} w_{i,d} \log f\left(r_{i,d}^{(\tau)} \mid \beta_{0,\tau}, \beta_{A2A,\tau}, \beta_{vol,\tau}, \beta_{liq,\tau}; \sigma^2\right),$$

where $f(\cdot \mid \cdot)$ denotes the conditional density implied by the GLM family and link function. In this formulation, the coefficient $\beta_{A2A,\tau}$ continues to capture the conditional mean difference in price impact between A2A and Voice trades, but under a design that balances the influence of each class on a daily basis.

It is important to stress that, given the relative stability of class proportions in our sample, a standard OLS or unweighted GLM would yield similar estimates. Our motivation for daily weighting lies instead in improving inference precision. In particular, the variance of estimated coefficients depends on the curvature of the log-likelihood, and weighting helps mitigate distortions in confidence interval estimation that could otherwise arise from uneven daily class representation.

4 Results

We estimate an *impact curve* by calculating the coefficients $\beta_{A2A,\tau}$ at different time horizons τ , with steps of five minutes up to 100 minutes after an event. Importantly, all estimates are obtained prior to the trade being reported to TRACE, ensuring that the results capture pre-trade information effects of A2A execution rather than post-trade market adjustments.

For interpretability, we align the direction of the outcomes so that a positive coefficient corresponds to a price move in the same direction as the trade. In other words, for SELL trades we expect spreads to rise (widen), while for buy trades we expect spreads to fall (tighten). A positive $\beta_{A2A,\tau}$ therefore indicates that A2A execution is associated with larger same-direction price impact relative to Voice trades.

Figures 1 and 2 summarize the estimated conditional effects. The y-axis reports the mean differential in spreads (in basis points), conditional on the trade direction and covariates. Across both buy and sell segments, we observe significant same-direction price pressure following A2A executions. The magnitude of the effect is systematically larger for sell-initiated trades, which display a more persistent widening of spreads relative to Voice Trades. In contrast, buy-initiated trades exhibit a positive but comparatively attenuated tightening effect.

Liquidity Segmentation

To further investigate the heterogeneity of the A2A impact across liquidity regimes, we conduct a segmentation analysis based on the previously defined liquidity score. Specifically, we re-estimate the model separately for bonds falling within the bottom 50% (low-liquidity) and top 50% (high-liquidity) of the 90-day rolling liquidity score distribution. This split allows for a direct comparison of A2A effects conditional on market depth and trading activity.

The results, displayed in Figures 3 and 4, reveal that the magnitude and persistence of the A2A effect depend strongly on liquidity conditions. In the most liquid segment, A2A executions exhibit relatively muted same-direction price adjustments, consistent with tighter spreads, faster information diffusion, and higher market resiliency. Conversely, in the least liquid segment, A2A trades are associated with significantly larger and more persistent same-direction price movements, suggesting that automation may amplify price impact under constrained liquidity. These patterns are consistent with market microstructure theory, where limited liquidity increases the sensitivity of prices to trading imbalances.

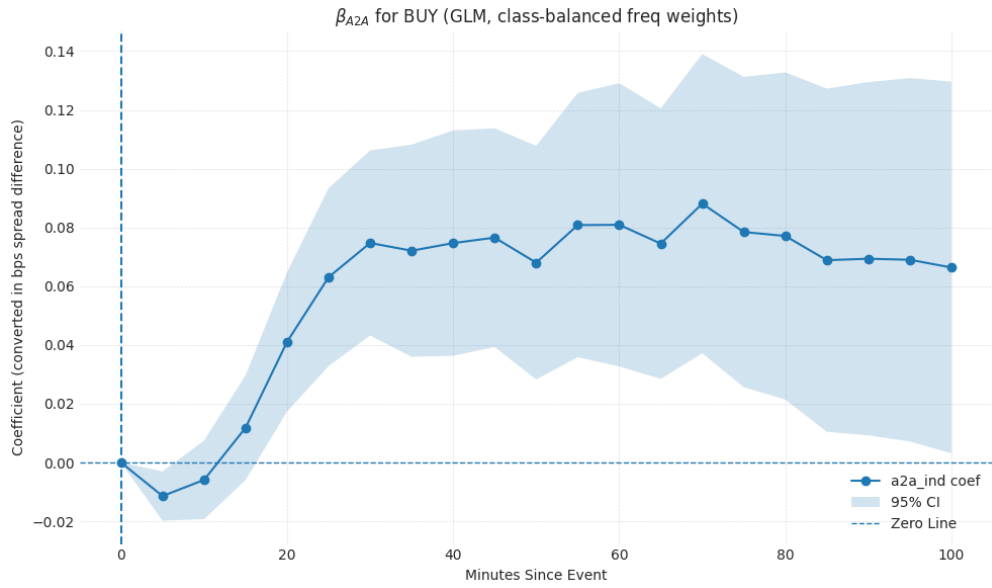


Figure 1: Estimated impact curve for buy-initiated trades. The y-axis represents the conditional mean difference in spreads (basis points) between A2A and Voice executions.

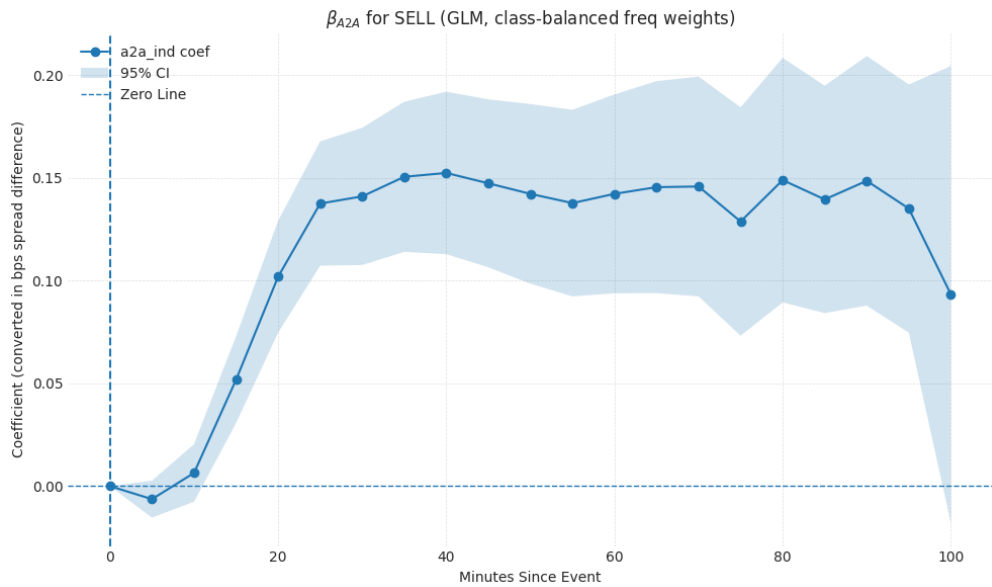


Figure 2: Estimated impact curve for sell-initiated trades. The y-axis represents the conditional mean difference in spreads (basis points) between A2A and Voice executions.

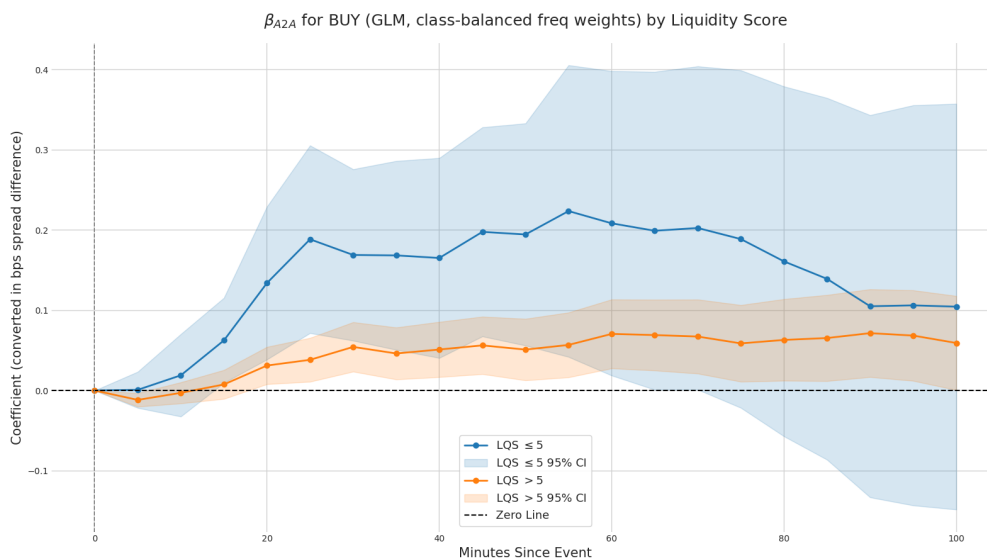


Figure 3: Impact curves for buy-initiated trades, segmented by liquidity. The figure reports the estimated conditional mean difference in spreads (basis points) between A2A and Voice executions, computed separately for bonds in the top 50% (high-liquidity) and bottom 50% (low-liquidity) of the 90-day rolling liquidity score distribution. A positive value indicates a larger same-direction price movement for A2A executions relative to Voice trades.

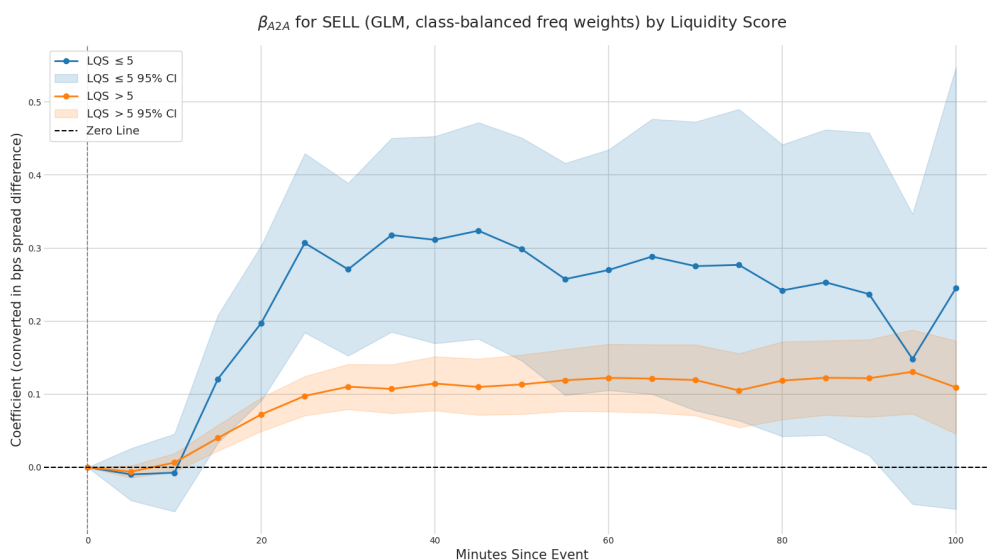


Figure 4: Impact curves for sell-initiated trades, segmented by liquidity. The y-axis shows the conditional mean difference in spreads (basis points) between A2A and Voice executions, estimated separately for high- and low-liquidity bonds (top and bottom 50% of the 90-day rolling liquidity score distribution). The results illustrate how liquidity conditions influence the magnitude and persistence of same-direction price pressure associated with A2A execution.

5 Discussion and Conclusion

This study examines the differential impact of execution protocols - specifically, A2A versus Voice - on subsequent price dynamics in the corporate bond market. By employing a generalized linear framework with extensive controls such as trade volume and liquidity, we sought to orthogonalize the protocol effect from known confounding influences. Our empirical design isolates the conditional response of spreads to trade direction and execution channel, ensuring that the estimated A2A coefficients reflect information effects and inventory adjustments intrinsic to the execution mechanism rather than broader market factors.

Our identification strategy requires the assumption that after conditioning on the included covariates, the assignment of trades to A2A or Voice protocols is conditionally independent of unobserved drivers of subsequent price movements. However, this assumption is partially justified by the high-frequency nature of our estimation horizon and the inclusion of controls that capture primary dimensions of market microstructure behavior influencing transaction costs and price pressure. Within this framework, the estimated A2A effect could be interpreted as quasi-causal, reflecting the incremental contribution of A2A trading to short-term price adjustments.

From a market perspective, the findings point to a directionally consistent price impact following A2A executions, quantifiable for buy- and sell-initiated trades in the order of 0.1 to 0.3 basis points. We interpret these results as an emerging and increasing information channel that will likely become increasingly relevant as the market pushes for more electronification and transparency.

The implications of these results differ across participants:

- Buy-side institutions benefit from improved access and transparency in A2A protocols but must remain aware of the potential adverse selection effects associated with large trades - particularly those exceeding one million in notional value. Strategic timing or selective disclosure might mitigate these mechanisms.
- Dealers, conversely, are increasingly exposed to adverse price action following A2A executions on the opposite side of the trade.

In conclusion, our findings highlight the complex and evolving nature of information transmission in modern bond trading. While the A2A trading protocol promotes market transparency, liquidity and competition, the shift in the distribution of information and price discovery dynamics has also become increasingly noticeable. Continuing to monitor these dynamics will be crucial as the markets adapt and converge towards increasingly data-driven and automated architectures. These dynamics support the development of hybrid execution approaches. Tradeweb's SNAP+ protocol enables more targeted dealer selection based on historical performance and real-time activity, preserving competition while limiting information exposure. This approach offers a middle ground between bilateral Voice execution and fully distributed A2A trading.

References

- Bouchaud, J.P. (2009). Price Impact. arXiv:0903.2428 [q-fin]. [online] Available at: <https://arxiv.org/abs/0903.2428>.
- Guo, X., Lehalle, C.-A. and Xu, R. (2019). Transaction Cost Analytics for Corporate Bonds. [online] arXiv.org. Available at: <https://arxiv.org/abs/1903.09140> [Accessed 12 May 2026].
- HASBROUCK, J. (2009). Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data. *The Journal of Finance*, 64(3), pp.1445–1477. doi:<https://doi.org/10.1111/j.1540-6261.2009.01469.x>.
- O’Hara, M. and Zhou, X. (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics*, 140(2), pp.368-390. doi:<https://doi.org/10.1016/j.jfineco.2021.01.001>.
- Demina, A., Mahajan, H. and Husveth, T. (2026) “From ”Blast All” to ”Select Smart”: How Dealer Selection Is Evolving.” Tradeweb Insights. <https://www.tradeweb.com/newsroom/media-center/insights/blog/from-blast-all-to-select-smart-how-dealer-selection-is-evolving>.

Author Contacts

Filippo Caretti +1 617 233 8031
Assistant Vice President, U.S. Credit
filippo.caretti@tradeweb.com

Harveer Mahajan +1 617 852 9473
Director, U.S. Credit
harveer.mahajan@tradeweb.com

Ted Husveth +1 646 560 7555
Managing Director, U.S. Credit
ted.husveth@tradeweb.com

Media Relations

Daniel Noonan +1 646 767 4677
daniel.noonan@tradeweb.com

Savannah Steele +1 646 767 4941
savannah.steele@tradeweb.com