ANNUAL REVIEWS

ANNUAL REVIEWS CONNECT
www.annualreviews.org
• Download figures
• Navigate cited references
• Keyword search
• Explore related articles
• Share via email or social media

t (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

Annual Review of Financial Economics

The Evolution of the Corporate Bond Market: A Theoretical Analysis

Mahyar Kargar,[1] Benjamin Lester,[2] Semih Üslü,[3,4] and Pierre-Olivier Weill[5,6,7]

[1]Gies College of Business, University of Illinois Urbana-Champaign, Champaign, Illinois, USA; email: kargar@illinois.edu

[2]Federal Reserve Bank of Philadelphia, Philadelphia, Pennsylvania, USA; email: benjamin.lester@phil.frb.org

[3]Desautels Faculty of Management, McGill University, Montreal, Quebec, Canada; email: semih.uslu@mcgill.ca

[4]Carey Business School, Johns Hopkins University, Baltimore, Maryland, USA

[5]Department of Economics, University of California, Los Angeles, California, USA; email: poweill@econ.ucla.edu

[6]National Bureau of Economic Research, Cambridge, Massachusetts, USA

[7]Centre for Economic Policy Research, London, United Kingdom

Annu. Rev. Financ. Econ. 2026. 18:161–84

First published as a Review in Advance on August 28, 2026

The *Annual Review of Financial Economics* is online at financial.annualreviews.org

https://doi.org/10.1146/annurev-financial-111424-014232

Copyright © 2026 by the author(s). This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0), which permits any noncommercial use, sharing, distribution, and reproduction in any medium or format, provided the original author(s) and source are credited; this license does not permit sharing adapted material derived from this article or parts of it. Images or other third-party material in this article are included in the article's Creative Commons license unless indicated otherwise; see credit lines for license information.

JEL codes: G11, G12, G14, G24, D83

OPEN ACCESS

Keywords

over-the-counter markets, corporate bond market liquidity, dealer intermediation, balance sheet costs, electronic trading platforms

Abstract

We develop a model of a dealer-intermediated over-the-counter market designed to study three major changes in the structure of the US corporate bond market: the increase in dealers' balance sheet costs, the emergence of electronic trading platforms, and the growing presence of bond mutual funds and exchange-traded funds. Our model provides a unified analysis of these changes, clarifies the economic channels at play, and allows us to quantify their effects on a variety of market outcomes. Our quantitative analysis suggests that while electronic trading significantly reduced the cost of raising capital in the corporate bond market, these gains were almost completely offset by the combined effects of balance sheet costs and changes in the demand for liquidity. We find that electronic trading also caused a meaningful decline in the bid-ask spread, whereas other changes in the market structure had little effect on transaction costs.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

1. INTRODUCTION

The corporate bond market, the primary source of long-term borrowing for large American companies, has changed in several significant ways over the past two decades.[1] First, after the global financial crisis (GFC), many dealers in the corporate bond market were subject to regulations that increased the costs of holding assets on their balance sheet. Since dealers' "balance sheet space" is a crucial component of their ability to intermediate trades, many have argued that these regulations inadvertently made dealers less willing to supply liquidity to investors. Second, a sizable share of trading volume has migrated from traditional voice-based trades—i.e., bilateral transactions that take place over the phone or via chat—to electronic trading platforms that allow investors to contact multiple dealers at once, making it easier to locate a particular bond and fostering competition among dealers.[2] Third, traditional owners of corporate bonds, such as life insurance companies and pension funds, have been gradually replaced by bond mutual funds and exchange-traded funds (ETFs), which have starkly different trading needs.[3]

A recent literature has studied the effects of these shifts on trading activity, transaction costs, and corporate bond yield spreads. This work has been predominantly empirical, typically analyzing events in isolation to establish causal effects. We provide a complementary perspective by developing a theoretical model of a dealer-intermediated over-the-counter (OTC) market, allowing for a unified analysis of the changes in the corporate bond market discussed above, the economic channels at play, and their quantitative impact on market outcomes.

We consider a continuous-time model of an OTC market for a representative bond. The market is populated by two types of agents—customers and dealers—both of whom can hold either zero or one unit of the asset. Customers have time-varying and heterogeneous preferences for the asset, which generates gains from trade. At stochastic arrival times, these customers have the opportunity to trade through dealers, who have stable, homogeneous preferences for the asset. Dealers, however, have access to a frictionless interdealer market where they can buy and sell the asset at a competitive price.

To this baseline model, we add two key ingredients. First, we assume that customers and dealers come into contact in two ways: through voice meetings, in which a customer contacts a single dealer and they bargain over the price if there are gains from trade, and via an electronic trading platform, in which the customer contacts multiple dealers, each of whom may respond with an offer. The second key ingredient is that balance sheet space serves as an essential input for dealers' intermediation activity. When contacted by a customer, a dealer can sell the asset only if they currently own it, and they can buy the asset only if they have available balance sheet space (i.e., they do not already own the asset).

We calibrate the model to match empirical moments from the data, and conduct comparative statics to explore the three changes in the corporate bond market discussed above. While the calibrated model yields a variety of testable implications, we focus our attention on understanding the behavior of—and the connection between—two important market outcomes: the bid-ask spread, often considered a proxy for the quality of secondary asset markets, and the liquidity yield spread,

[1] According to the Securities Industry and Financial Markets Association (SIFMA), as of October 2025, outstanding corporate debt exceeded $11 trillion while average daily trading volume in the corporate bond market was nearly $60 billion (**https://www.sifma.org/research/statistics/us-corporate-bonds-statistics**).

[2] As we describe in greater detail in Section 2, the share of electronic trading increased dramatically in the last decade. In some segments of the market—specifically, smaller trades for IG bonds—the share of volume that has migrated to electronic trading platforms is more than half.

[3] Table L.213 of the Flow of Funds (Federal Reserve System 2025) shows that from 2000:Q1 to 2025:Q2, the share of corporate and foreign bonds held by mutual funds and ETFs surged from 9.5% to 23% while the share held by life insurance companies and pension funds decreased from 41% to 32%.

which affects the real cost of capital, firms' investment, and hence, the macroeconomic effects of secondary market frictions (as formalized by Kozlowski 2021).[4]

The first takeaway of our numerical exercise is that while electronic trading significantly reduced the cost of raising capital in the corporate bond market, these gains were almost completely offset by the combined effects of changes in balance sheet costs and customers' trading needs. We find that electronic trading also caused a meaningful decline in the bid-ask spread, whereas other changes in the market structure had little effect on transaction costs. Hence, a second takeaway is that changes in market structure do not necessarily move the yield spread and bid-ask spread in unison.

To understand these findings, let us consider each change in market structure in isolation. We find that a rise in dealers' balance sheet costs leads to a sizable increase in the yield spread but has almost no effect on bid-ask spreads, which is consistent with empirical studies that have struggled to measure meaningful changes in the bid-ask spread in response to the implementation of post-GFC regulations (e.g., Bessembinder et al. 2018, Trebbi & Xiao 2019). A similar pattern emerges when we increase the frequency of customers' trading needs, meant to reflect the change in the composition of investors: We find that the liquidity yield spread rises while bid-ask spreads remain roughly unchanged, which is consistent with the empirical findings of Li & Yu (2026). However, we find that an increase in the prevalence of electronic trading creates a sizable reduction in bid-ask spreads and yield spreads. Putting it all together, when we study the total impact of all three changes in market structure, we find that the effects on the liquidity yield spread essentially cancel each other out, while the decline in the bid-ask spread remains economically significant.

Our model reveals that the relationship between these two market outcomes can be understood through the condition that determines the dealers' optimal use of balance sheet space. In particular, dealers can utilize their balance sheet to carry inventory sourced from the interdealer market—incurring exogenous holding costs to service ask-side demand—or they can reserve capacity to accommodate bid-side supply from customers. In equilibrium, dealers are active on both sides of the market and hence must be indifferent between these two choices. This optimality condition links three key components of dealers' profits. The first component is the ask-side profits, i.e., the expected value of acquiring the asset and selling it to a customer. The second component is the liquidity yield spread, an indirect monetary benefit from holding the asset. And the third component is the bid-side profits, i.e., the expected value of keeping balance sheet space in order to buy an asset from a customer. The dealers' optimality condition requires that the sum of ask-side profits and the liquidity yield spread is equal to the sum of bid-side profits and exogenous balance sheet costs.

This condition delivers two key insights. The first key insight is that the relationship between the level of transaction costs and liquidity yield spreads depends on two new empirical measures—the ask-side and bid-side profits per unit of bond—which are calculated as the product of the order flow per dealer with the ask and bid spreads, respectively. These measures can be calculated with information on bid spreads, ask spreads, bond turnover, and the size of the dealer sector. The second key insight is that any change in market structure leads to compensating changes in ask-side profits, bid-side profits, and the liquidity yield spread. To analyze these compensating changes, of course, one needs to go beyond the analysis of optimality conditions and analyze the full OTC market equilibrium. We thus establish equilibrium existence, provide conditions for uniqueness, and offer numerical comparative statics: We measure these compensating changes in a calibrated version of our model designed to match empirical moments from the corporate bond market.

[4]Since we study an environment with risk-free bonds, our model cannot speak to the impact of liquidity on the credit component of the yield spread (as formalized by He & Milbradt 2014).

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

To begin, we consider an increase in dealers' balance sheet costs of 32 basis points (bps), as estimated by Fleckenstein & Longstaff (2020) in the context of the Treasury market. When balance sheet costs rise, dealers have less incentive to hold the asset; we find that in equilibrium, the liquidity yield spread rises by approximately 5 bps, and the fraction of bonds held in the dealer sector falls by approximately 6%. This means that there is less competition among dealers who own assets, such that ask prices rise and ask-side profits increase by approximately 11 bps. At the same time, since more dealers have balance sheet space, bid prices become more competitive and bid-side profits fall by approximately 16 bps. With both bid and ask prices rising, we find that the bid-ask spread is little changed. Hence, a key takeaway of this numerical example is that bid and ask spreads need not change much to restore dealers' incentives to provide liquidity in the face of an increase in balance sheet costs, which is consistent with empirical studies.

Next, we consider an increase in the rate at which customers can submit a request for quotes (RFQ) on an electronic platform, such that our model generates an increase in the share of electronic trading from approximately 20% to 50%, consistent with empirical estimates for investment-grade (IG) bonds (Coalition Greenwich 2015; McPartland 2023, 2025). We find that bid- and ask-side profits decline by much larger amounts, generating a sizable impact on the bid-ask spread. Finally, we consider an increase in the frequency of customers' trading needs, intended to capture the shift from traditional buy-and-hold investors toward investment vehicles that rebalance their portfolios more often. As noted above, we find that the liquidity yield spread rises while bid-ask spreads remain roughly unchanged, as in the data.

We conclude this section with a brief note on the literature most closely related to our work. Our market setting is a variation of the canonical search-theoretic model of OTC markets developed by Duffie, Gârleanu & Pedersen (2005); we refer the reader to Hugonnier, Lester & Weill (2025) for a comprehensive overview of this literature and its methods. As noted above, we introduce two key ingredients, relative to the standard model.

To start, building on Glebkin, Yueshen & Shen (2023), we incorporate multilateral RFQs to study the emergence of electronic trading platforms, though we extend their theoretical market setting in two important ways. First, we introduce a competitive interdealer market to study the determination of an equilibrium interdealer price, which is conceptually similar to the reference price calculated in empirical studies (see, e.g., Choi, Huh & Shin 2024). This facilitates the analysis of the liquidity yield spread and the bid- and ask-side profits and spreads. Second, we extend their two-type model to the case of a continuum of customer types, following Hugonnier, Lester & Weill (2022). With two types, in their preferred region of the parameter space, only the equilibrium price adjusts to a marginal change in balance sheet costs. In our setting with a continuum of types, quantities adjust as well; for instance, when balance sheet costs increase, the share of bonds held by dealers falls, which is consistent with empirical observations.

The second key ingredient, borrowed from Cohen et al. (2024), is the assumption that balance sheet space is a key input for intermediation activity. Importantly, this assumption generates an incentive for dealers to hold assets on their balance sheet even though there is a competitive interdealer market. Relative to Cohen et al. (2024), we make the simplifying assumption that customers and dealers can only hold zero or one unit of the asset, but crucially, we allow customers to use two trading mechanisms, voice and electronic.

The article proceeds as follows. Section 2 offers a brief description of the three recent changes in the US corporate bond market that we have chosen to study. Section 3 introduces the model, characterizes the equilibrium, and explores some of the relationships between the economic environment and equilibrium outcomes. In Section 4, we calibrate the model and explore the quantitative effects of these changes in the market structure. Section 5 discusses some of the limitations of our model framework, potential extensions and modifications to address these

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

limitations, and avenues for future research. Proofs of all propositions are provided in the **Supplemental Appendix**.

Supplemental Material >

2. BACKGROUND

In this section, we briefly describe the three major changes in the corporate bond market that we study both theoretically and quantitatively. For more detailed discussions of these changes—and other recent developments—we refer the reader to the comprehensive review of the corporate bond market by O'Hara & Zhou (2025).

2.1. Regulation and Balance Sheet Costs

The early 2010s was a period of significant regulatory reform, with a variety of stricter requirements placed on depository institutions in the wake of the GFC. For example, the 2010 Basel III framework introduced enhanced capital and liquidity requirements on banks, while the Volcker rule introduced greater scrutiny over banks' inventory holdings in order to enforce the rule's ban on proprietary trading. While these regulations reduced banks' leverage and arguably made the banking sector as a whole less fragile, many argued that they also increased the cost for bank-affiliated dealers of holding assets on their balance sheets and hence reduced their incentives to provide intermediation services.[5]

Following the implementation of these regulatory reforms, there is evidence that affected dealers curtailed their market-making activities and that nonaffected dealers did not fully offset this contraction. For example, the dealer sector as a whole reduced the level of assets on their balance sheets: According to Table L.213 of the Flow of Funds (Federal Reserve System 2025), the share of outstanding corporate bonds held by security brokers and dealers fell from 2–3% in 2006 to less than 1% in 2018.[6] Studying differences across dealers and over time, a variety of empirical studies have documented that post-GFC regulations reduced affected dealers' capital commitment and increased price impact during episodes of high selling pressure, though there is less consensus on the ultimate effects on bid-ask spreads (see, e.g., Bao, O'Hara & Zhou 2018; Bessembinder et al. 2018; Trebbi & Xiao 2019; Dick-Nielsen & Rossi 2019; Choi, Huh & Shin 2024).

Highlighting the importance of dealers' balance sheet space in providing liquidity, several studies have documented that bond-level liquidity comoves with intermediaries' balance sheet tightness: When constraints bind, transaction costs rise and dealers' willingness to take inventory falls, consistent with a state-dependent "price of immediacy" (see, e.g., Adrian, Boyarchenko & Shachar 2017; Lou et al. 2025). Liquidity conditions during the COVID-19 episode provided further evidence that dealers had become less willing to "lean against the wind"; for example, Kargar et al. (2021) document that the cost of risky principal trades (where dealers absorb bonds into their balance sheets) spiked in March 2020 and trading migrated to agency trades (where dealers match customer buyers and customer sellers without holding the asset themselves).

2.2. The Rise of Electronic Trading Platforms

Electronic trading first emerged in the corporate bond market in the 1990s and 2000s, as companies like MarketAxess and Tradeweb offered investors the opportunity to contact multiple dealers

[5]For a more detailed discussion of the effects of post-GFC regulations on liquidity provision, readers are referred to Duffie (2012), Thakor (2012), Duffie (2017), Bessembinder et al. (2018), and the many references therein.

[6]The data also include foreign bonds, privately issued mortgage-backed securities, and other privately issued asset-backed bonds.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

at once via RFQs in electronic platforms. The market share of these electronic platforms has grown dramatically: According to Coalition Greenwich (2015) and McPartland (2025), the share of electronic trading for IG bonds grew from approximately 20% in 2015 to approximately 50% in 2025. As electronic trading gained market share, O'Hara & Zhou (2021) document that transaction costs fell for both electronic and voice trades. It is important to note, however, that voice-based trading remains a significant fraction of trading volume, particularly in certain segments of the corporate bond market. In particular, electronic trading is most pronounced for smaller, IG trades, while large blocks and lower-rated issues remain predominantly negotiated by voice (see, e.g., Hendershott & Madhavan 2015, O'Hara & Zhou 2021, Kargar et al. 2025).

Yet even within the segment of the market in which electronic trading has grown rapidly, customers are not always able to get good quotes by submitting an RFQ. Kargar et al. (2025) establish that inventory is an important determinant of dealers' willingness and ability to provide quotes. In particular, they establish that dealers with a relatively large inventory of a particular bond are significantly more (less) likely to respond to a customer-buy (customer-sell) RFQ and, conditional on replying, quote a significantly better ask price (worse bid price).

2.3. The Rise of Bond Mutual Funds and ETFs

In the nearly two decades since the GFC, the composition of investors who own corporate bonds has shifted from buy-and-hold investors like insurers and pension funds to open-ended bond mutual funds and ETFs that actively respond to changing market conditions. For example, according to Table L.213 of the Flow of Funds (Federal Reserve System 2025), the share of outstanding corporate bonds held by mutual funds and ETFs increased from less than 10% before the GFC to more than 20% by the second quarter of 2025.[7]

This change in the composition of corporate bond owners has a variety of potential effects on the demand for liquidity. The literature following Goldstein, Jiang & Ng (2017) has emphasized that the rise of open-ended funds and ETFs has increased the sensitivity and magnitude of liquidity demand shocks, as outflows from these investment vehicles respond strongly to poor performance and aggregate macroeconomic conditions. In particular, recent empirical studies have documented that bond mutual funds trade more actively and are more price elastic than traditional buy-and-hold investors (Timmer 2018, Bretscher et al. 2026), while ETFs tend to attract investors with greater liquidity demand (Manconi, Massa & Yasuda 2012; Dannhauser & Hoseinzade 2022).

2.4. Summary

Taken together, these facts motivate the main ingredients that we incorporate into our theoretical framework. First, we model an intermediation process in which dealers trade off the costs of holding inventory with the benefits of providing liquidity to customers who are searching to buy. Second, we allow for two trading technologies, intended to capture voice trades and electronic RFQs. Lastly, we introduce a preference parameter that determines the frequency of customers' need to trade.

3. MODEL

In this section, we describe our theoretical framework, derive the optimal behavior of customers and dealers, and characterize equilibrium prices and allocations. We then explore the relationship between several model parameters—designed to capture key features of the market structure—and observable outcomes, including bid prices, ask prices, and the liquidity yield spread.

[7] For more details, readers are referred to Li & Yu (2026).

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

3.1. Setup

We consider an economy populated by two types of infinitely lived, risk-neutral agents who discount future payoffs at rate $r > 0$: a measure 1 of customers and a measure m of dealers. All agents can hold either zero or one unit of a perpetual bond in supply $s \in (m, m + 1)$. We let the bond coupon rate be $r + \ell$, where the spread ℓ will be determined in equilibrium so that dealers trade the bond at a par value of $P = 1$.[8] This means that in our model, ℓ is equal to the yield spread, i.e., the adjustment to the discount rate required for the price to equal the present value of the coupon payments. We refer to ℓ as the "liquidity yield spread," since in our framework the yield spread derives only from liquidity considerations (as opposed to, e.g., risk).

As is standard in the literature, customers need to trade because their preferences for the bond vary stochastically over time. Specifically, we assume that a customer who holds the bond at a given time derives a total utility flow $r + \ell + \varepsilon$, where ε is a convenience yield. We allow the convenience yield to take on negative values, which we interpret as a distress cost. Each customer draws a new convenience yield at independent Poisson arrival times with intensity γ, according to the continuous cumulative distribution function F with support $[\underline{\varepsilon}, \overline{\varepsilon}]$.

Dealers can hold zero or one unit of the asset as well, in which case they derive a time-invariant flow utility $r + \ell + \varepsilon_D$. The convenience yield ε_D represents the liquidity benefits of holding the assets net of balance sheet costs that dealers incur. Later, we use a decrease in ε_D to represent the increase in balance sheet costs that dealers incurred as a result of the regulations put in place after the GFC.

The OTC market for the bond is semicentralized: Customers must search for dealers in order to trade, while dealers can trade together in a frictionless competitive market. But we depart from the standard semicentralized model described in Hugonnier, Lester & Weill (2025, chapter 4) by incorporating model elements from Glebkin, Yueshen & Shen (2023) and Cohen et al. (2024).

First, we assume that customers and dealers come into contact through one of two different technologies: With intensity λ_v customers contact dealers via a voice technology, and with intensity λ_e customers contact dealers via an electronic trading platform. In the former case, the investor contacts a single dealer and they bargain over the price if there are gains from trade. In the latter case, meant to capture an electronic RFQ, the investor contacts $N > 1$ dealers, each of whom may respond with an offer.[9] Importantly, when a dealer who receives an electronic RFQ responds, they do not know how many other dealers will also respond, as in Burdett & Judd (1983).

Second, we assume that a dealer must choose their inventory holdings, which determines the direction of trade should they come into contact with a customer. A dealer who holds a unit of the asset can sell to a customer but cannot buy, while a dealer who does not own the asset—i.e., who has balance sheet space—can buy but not sell. This assumption implies that dealers face a trade-off between providing liquidity on the bid side (to customer sellers) and the ask side (to customer buyers).

3.2. Bid and Ask Prices in Voice and Electronic Trades

In this subsection, we describe how prices are determined in voice and electronic trades and derive expressions for these prices. To do so, let $V(q, \varepsilon)$ denote the maximum expected utility of a customer

[8] Assuming that the bond trades at par is a normalization, since the model is linearly homogeneous: Equilibrium quantities remain unchanged if one scales up or down the utility flows and the prices by the same constant.

[9] Taking N as given is a natural assumption because, in practice, most electronic trading platforms assign to each customer a set of dealers who receive their RFQs, often based on preexisting trading relationships. We discuss this at greater length in Section 5, along with the implications of allowing customers to choose the number of dealers, N, who receive an RFQ.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

with asset holding $q \in \{0, 1\}$ and convenience yield ε. As is standard, a customer's reservation value for holding a unit of the asset is $R(\varepsilon) = V(1, \varepsilon) - V(0, \varepsilon)$, the difference between the value of owning and not owning a unit of the asset, when trading optimally. Since dealers can trade instantly with each other before and after meeting customers, their reservation value is equal to the price of the asset in the perfectly competitive interdealer market, P.

3.2.1. Voice trades. Consider a bilateral meeting via voice between a dealer who owns one unit of the asset and a customer nonowner with convenience yield ε. Clearly, there are gains from trade if and only if $R(\varepsilon) > P$. Consistent with the literature, we assume that the ask price is determined by the generalized Nash bargaining solution,

$$\theta_0 R(\varepsilon) + (1 - \theta_0)P,$$

where θ_q denotes the bargaining power of a dealer when trading with a customer who owns $q \in \{0, 1\}$ unit of the asset. Symmetrically, in a meeting between a dealer nonowner and a customer owner with convenience yield ε, there are gains from trade if and only if $R(\varepsilon) < P$, and the bid price is

$$\theta_1 R(\varepsilon) + (1 - \theta_1)P.$$

3.2.2. Electronic trades. Consider now an electronic multilateral meeting, in which a customer nonowner with reservation value $R(\varepsilon) > P$ contacts N dealers. Since, in practice, these multilateral meetings are not anonymous, we make the simplifying assumption that dealers observe the convenience yield ε of the customer. As in Glebkin, Yueshen & Shen (2023), although customers contact a fixed number of dealers, N, they receive a random number of quotes, determined by the number of contacted dealers who currently hold the asset. Specifically, assuming RFQs are sent to dealers randomly, the number of dealers who can respond with a quote is distributed according to a binomial distribution with parameter $\phi_1 = 1 - \phi_0$, where ϕ_q denotes the fraction of dealers who own $q \in \{0, 1\}$ unit of the asset.

It is well-known from Burdett & Judd (1983) that in the equilibrium of this price setting game, dealers play mixed strategies: They choose a price at random according to a continuous distribution that is characterized analytically in Hugonnier, Lester & Weill (2025, chapter 8). The lowest quoted price is strictly greater than the competitive price, P, but the highest price is equal to the fully discriminating monopoly price, $R(\varepsilon)$. The optimal dealer strategy implies that trade will always occur if at least one dealer responds with a quote. Moreover, the average transaction price is

$$\Theta(\phi_0)R(\varepsilon) + (1 - \Theta(\phi_0))P,$$

where

$$\Theta(\phi) \equiv \frac{N(1 - \phi)\phi^{N-1}}{1 - \phi^N}.$$

The function $\Theta(\phi)$ is easily shown to be increasing, with $\Theta(0) = 0$ and $\lim_{\phi \to 1} \Theta(\phi) = 1$. Indeed, when $\phi_0 = 0$, customers always receive quotes from $N > 1$ dealers. Hence, dealers know that they are competing à la Bertrand and so, in equilibrium, quote their reservation value, P. Conversely, when $\phi_0 \to 1$, dealers anticipate no competing quotes and thus respond with the fully discriminating monopoly price, $R(\varepsilon)$.

The case of a customer who owns an asset with reservation value $R(\varepsilon) < P$ is symmetric: The average transaction price is

$$\Theta(\phi_1)R(\varepsilon) + (1 - \Theta(\phi_1))P.$$

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

3.3. The Customer's Problem

Consider a customer nonowner with convenience yield ε. The Hamilton–Jacobi–Bellman (HJB) equation for this customer is

$$\begin{aligned} rV(0,\varepsilon) = \gamma\left(\int V(0,x)\mathrm{d}F(x) - V(0,\varepsilon)\right) \\ + \lambda_{\mathrm{v}}(1-\phi_0)\Big[V(1,\varepsilon) - V(0,\varepsilon) - \theta_0 R(\varepsilon) - (1-\theta_0)P\Big]^+ \\ + \lambda_{\mathrm{e}}(1-\phi_0^N)\Big[V(1,\varepsilon) - V(0,\varepsilon) - \Theta(\phi_0)R(\varepsilon) - (1-\Theta(\phi_0))P\Big]^+, \end{aligned}$$

where $[x]^+ \equiv \max\{0, x\}$. The first term on the right-hand side is the (net) flow utility a customer receives when she draws a new convenience yield. The second term represents the (net) flow utility from a voice trade: The customer contacts a dealer owner with intensity $\lambda_{\mathrm{v}}(1-\phi_0)$, and if there are gains from trade, she acquires the asset. Upon trading, the customer experiences a change in value $V(1,\varepsilon) - V(0,\varepsilon)$ and pays the price derived above, $\theta_0 R(\varepsilon) + (1-\theta_0)P$. Similarly, the third term is the (net) flow utility from an electronic trade. Simplifying the formula, we obtain

$$\begin{aligned} rV(0,\varepsilon) = \gamma\left(\int V(0,x)\mathrm{d}F(x) - V(0,\varepsilon)\right) \\ + \lambda_{\mathrm{v}}(1-\phi_0)(1-\theta_0)\Big[R(\varepsilon) - P\Big]^+ + \lambda_{\mathrm{e}}(1-\phi_0^N)(1-\Theta(\phi_0))\Big[R(\varepsilon) - P\Big]^+. \end{aligned} \tag{1.}$$

Equation 1 reveals that a customer's value function can be calculated "as if" she could trade directly in the interdealer market at price P but with a bargaining-adjusted search intensity—that is, a search intensity adjusted for her share of the surplus. Similarly, the value function of a customer owner is

$$\begin{aligned} rV(1,\varepsilon) = r + \ell + \varepsilon + \gamma\left(\int V(1,x)\mathrm{d}F(x) - V(1,\varepsilon)\right) \\ + \lambda_{\mathrm{v}}(1-\phi_1)(1-\theta_1)\Big[P - R(\varepsilon)\Big]^+ + \lambda_{\mathrm{e}}(1-\phi_1^N)(1-\Theta(\phi_1))\Big[P - R(\varepsilon)\Big]^+. \end{aligned} \tag{2.}$$

Taking the difference between Equation 1 and Equation 2 yields the HJB for the reservation value function:

$$rR(\varepsilon) = r + \ell + \varepsilon + \gamma\left(\int R(x)\mathrm{d}F(x) - R(\varepsilon)\right) - \chi_0\Big[R(\varepsilon) - P\Big]^+ + \chi_1\Big[P - R(\varepsilon)\Big]^+, \tag{3.}$$

where

$$\chi_q \equiv \lambda_{\mathrm{v}}\left(1-\phi_q\right)\left(1-\theta_q\right) + \lambda_{\mathrm{e}}(1-\phi_q^N)\left(1-\Theta(\phi_q)\right)$$

denotes the trading intensities of customers holding $q \in \{0, 1\}$ unit of the asset, appropriately adjusted for the share of the surplus they appropriate in voice and electronic trades. Notably, in contrast to the standard semicentralized model, the bargaining-adjusted intensities are different on the ask side and the bid side since, in general, $\phi_1 \neq \phi_0$.

Note that the distribution of dealers' asset holdings affects customers' reservation values for two different reasons. First, a change in ϕ_q, for $q \in \{0, 1\}$, affects the probability that the customer finds a dealer with whom to trade; for example, an increase in ϕ_0 makes it less likely that a customer buyer will be matched with a dealer who owns a unit of the asset. Second, a change in ϕ_q also changes the level of competition among dealers who can make an offer; for example, an increase in ϕ_0 decreases competition among dealers who are responding to an RFQ from a customer buyer

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

and hence shifts the distribution toward less generous offers, leaving the customer buyer with a smaller fraction of the surplus.

Let us define $\varepsilon^\star$ as the convenience yield of a marginal customer, i.e., a customer with $R(\varepsilon^\star) = P$. Thus, customers buy the asset when $\varepsilon > \varepsilon^\star$, and they sell if $\varepsilon < \varepsilon^\star$. It follows from results in Hugonnier, Lester & Weill (2025) that the reservation value function is strictly increasing in ε, differentiable everywhere except at $\varepsilon^\star$, with derivative

$$R'(\varepsilon) = \frac{1}{r + \gamma + \chi_0 \mathbb{I}_{\{\varepsilon > \varepsilon^\star\}} + \chi_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}}}. \tag{4.}$$

This expression for the derivative will prove useful later to derive explicit formulae for the liquidity yield spread and the bid-ask spread.

3.4. Distributions and Market Clearing

We now characterize the steady-state distribution of customers across asset holdings and preference types and derive the condition that clears the asset market. Together, these objects determine how the marginal customer, $\varepsilon^\star$, is related to the fraction of dealers who hold the asset, ϕ_1.

3.4.1. The joint distribution of asset holdings and preference types. Let $\Psi_q(\varepsilon)$ denote the cumulative measure of customers with convenience yield less than ε and asset holding $q \in \{0, 1\}$. Since the total measure of customers is normalized to 1, we must have

$$\Psi_0(\varepsilon) + \Psi_1(\varepsilon) = F(\varepsilon). \tag{5.}$$

Let

$$\lambda_q \equiv \lambda_v(1 - \phi_q) + \lambda_e(1 - \phi_q^N)$$

denote the intensities with which customers with asset holding $q \in \{0, 1\}$ meet dealers with whom they can feasibly trade. From the definition of $\varepsilon^\star$, a customer nonowner buys from a dealer if $\varepsilon > \varepsilon^\star$, while a customer owner sells to a dealer if $\varepsilon < \varepsilon^\star$.[10] Hence, in a steady-state equilibrium, the inflow and outflow of customers in state (q, ε) must be equal, such that

$$\lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} d\Psi_0(\varepsilon) + \gamma\,(s - m\phi_1)\, dF(\varepsilon) = \lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}} d\Psi_1(\varepsilon) + \gamma d\Psi_1(\varepsilon). \tag{6.}$$

The left-hand side of Equation 6 represents the inflow of customers into the set of owners with a convenience yield in $[\varepsilon, \varepsilon + d\varepsilon]$. First, at arrival rate λ_0, customer buyers in the set of nonowners with convenience yield in $[\varepsilon, \varepsilon + d\varepsilon]$ have the opportunity to trade, which is successful if $\varepsilon \geq \varepsilon^\star$. Second, among the measure $s - m\phi_1$ of customers who own the asset, a fraction γ will draw a new convenience yield ε with probability $dF(\varepsilon)$. In contrast, the right-hand side of Equation 6 represents the outflow of customers from the set of owners with convenience yield in $[\varepsilon, \varepsilon + d\varepsilon]$. Again, the first term represents the outflow due to trade, and the second term represents the outflow due to customers drawing new convenience yields. Since $d\Psi_0(\varepsilon) = dF(\varepsilon) - d\Psi_1(\varepsilon)$ from Equation 5, we obtain

$$d\Psi_0(\varepsilon) = \frac{\lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}} + \gamma(1 - s + m\phi_1)}{\gamma + \lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}}}\, dF(\varepsilon), \tag{7.}$$

$$d\Psi_1(\varepsilon) = \frac{\lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \gamma\,(s - m\phi_1)}{\gamma + \lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}}}\, dF(\varepsilon). \tag{8.}$$

[10] Customers with $\varepsilon = \varepsilon^\star$ are indifferent but they are in measure zero, so their behavior does not matter for equilibrium—we adopt the convention that they buy if they do not own and do not sell if they own.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

3.4.2. Market clearing. Next, we formulate the market-clearing condition requiring that the measure of customers and dealers holding a unit of the asset must equal the supply:

$$\int_{-\infty}^{+\infty} \mathrm{d}\Psi_1(\varepsilon) + m\phi_1 = s.$$

Using the explicit expressions for λ_0 and λ_1, we obtain

$$\frac{\gamma\,(s - m\phi_1)}{\gamma + \lambda_1} F(\varepsilon^\star) + \frac{\gamma\,(s - m\phi_1) + \lambda_0}{\gamma + \lambda_0}\,(1 - F(\varepsilon^\star)) + m\phi_1 = s. \qquad 9.$$

Proposition 1 establishes that the market-clearing condition (Equation 9) implicitly defines a relationship between the marginal customer, $\varepsilon^\star$, and the fraction of dealers holding the asset, ϕ_1.

Proposition 1 (Market-clearing schedule). For all $\varepsilon^\star \in \mathbb{R}$, Equation 9 has a unique solution, $\phi_1 = \Phi_1(\varepsilon^\star)$, for some continuous and increasing function $\Phi_1(\cdot)$ such that $\Phi_1(\varepsilon^\star) = 0$ for all $\varepsilon^\star \leq \underline{\varepsilon}$, $\Phi_1(\varepsilon^\star) \in (0, 1)$ for all $\varepsilon^\star \in (\underline{\varepsilon}, \overline{\varepsilon})$, and $\Phi_1(\varepsilon^\star) = 1$ for all $\varepsilon^\star \geq \overline{\varepsilon}$.

The relationship between $\varepsilon^\star$ and ϕ_1 is monotonically increasing. Intuitively, as $\varepsilon^\star$ increases, the mass of customer nonowners with $\varepsilon \geq \varepsilon^\star$ (customer buyers) shrinks while the mass of customer owners with $\varepsilon < \varepsilon^\star$ (customer sellers) expands, resulting in a higher steady-state concentration of assets within the dealer sector. Note that the function $\Phi_1(\varepsilon^\star)$ is well-defined for all $\varepsilon^\star$, extending beyond the support of $F(\cdot)$. For example, if $\varepsilon^\star \leq \underline{\varepsilon}$, then all customers find it optimal to buy. The only way this can be consistent with the market-clearing and steady-state conditions is if buying is not feasible; i.e., if no dealers hold the asset ($\phi_1 = 0$).

3.5. The Dealer's Problem

We now turn to the dealer's optimal use of balance sheet space. At each instant, dealers choose whether to acquire a unit of the asset in the interdealer market. Acquiring the asset uses balance sheet space but allows the dealer to be active selling to customers on the ask side of the market. Alternatively, since not acquiring the asset preserves balance sheet space, it allows the dealer to be active buying from customers on the bid side of the market.

Let Π_q denote the expected value of (flow) profits for a dealer who owns $q \in \{0, 1\}$ unit of the asset. The difference between the value of acquiring the asset and not can be written

$$\Delta \equiv r + \ell + \varepsilon_{\mathrm{D}} - rP + \Pi_1 - \Pi_0 = \ell + \varepsilon_{\mathrm{D}} + \Pi_1 - \Pi_0,$$

where the equality follows from our normalizing assumption that the bond trades at par ($P = 1$) in the interdealer market. Naturally, dealers have greater incentive to acquire the asset if ask-side profits (Π_1) are large relative to bid-side profits (Π_0); if the liquidity yield spread (ℓ) is relatively large, since this increases the monetary return of holding the asset; or if their convenience yield (ε_{D}) is relatively large. To unpack the dealer's optimality condition, we derive explicit formulae for ℓ, Π_1, and Π_0 as functions of the marginal customer, $\varepsilon^\star$, and of the fraction of dealers who hold the asset, ϕ_1.

Starting with the liquidity yield spread, recall that $P = R(\varepsilon^\star)$, by definition, and that the bond trades at par ($P = 1$), by assumption. Using the reservation value (Equation 3) evaluated at $\varepsilon^\star$, we obtain that $\ell = \mathcal{L}(\varepsilon^\star \mid \phi_1)$, where

$$\mathcal{L}\,(\varepsilon^\star \mid \phi_1) \equiv -\varepsilon^\star - \frac{\gamma}{r + \gamma + \chi_0} \int_{\varepsilon^\star}^{+\infty} (\varepsilon - \varepsilon^\star)\,\mathrm{d}F(\varepsilon) + \frac{\gamma}{r + \gamma + \chi_1} \int_{-\infty}^{\varepsilon^\star} (\varepsilon^\star - \varepsilon)\,\mathrm{d}F(\varepsilon). \qquad 10.$$

The first component of the liquidity yield spread, $-\varepsilon^\star$, reflects the convenience yield of the marginal customer, which increases the price and therefore lowers the spread. The second

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

component arises because the marginal customer has a precautionary motive for purchasing the bond: She anticipates that if she draws a higher convenience yield, $\varepsilon > \varepsilon^\star$, she will want to purchase the bond but will not be able to do so immediately due to search frictions. The third component arises symmetrically because the marginal customer anticipates bearing distress costs in the future: If she draws a lower convenience yield, $\varepsilon < \varepsilon^\star$, she will want to sell the asset but will not be able to do so immediately.

Next, consider the flow profits on the ask side:

$$\begin{aligned}\Pi_1(\varepsilon^\star|\phi_1) &\equiv \left(\frac{\lambda_v}{m}\theta_0 + \frac{\lambda_e N}{m}\phi_0^{N-1}\right)\int_{\varepsilon^\star}^{\infty}(R(\varepsilon)-P)\,\mathrm{d}\Psi_0(\varepsilon)\\ &= \left(\frac{\lambda_v}{m}\theta_0 + \frac{\lambda_e N}{m}\phi_0^{N-1}\right)\gamma\,(1-s+m(1-\phi_0))\int_{\varepsilon^\star}^{\infty}\frac{\varepsilon-\varepsilon^\star}{(r+\gamma+\chi_0)(\gamma+\lambda_0)}\,\mathrm{d}F(\varepsilon).\end{aligned}$$

This equality is established by substituting the explicit formula for $\mathrm{d}\Psi_0(\varepsilon)$ from Equation 7 and applying Equation 4, noting that $R(\varepsilon) - P = \int_{\varepsilon^\star}^{\varepsilon} R'(x)\,\mathrm{d}x = (\varepsilon - \varepsilon^\star)/(r+\gamma+\chi_0)$. In the formula, $(\lambda_v/m)\theta_0$ is the bargaining-adjusted intensity of voice trading opportunities—that is, the product of the arrival rate of voice meetings at each dealer, λ_v/m, and the dealer's bargaining power, θ_0. Along the same vein, $(\lambda_e N/m)\phi_0^{N-1}$ is the bargaining-adjusted intensity of electronic trading opportunities for a dealer who owns the asset. Since customers contact N dealers at a time, an individual dealer receives electronic RFQs at a flow rate of $\lambda_e N/m$. To understand the bargaining adjustment, recall that dealers are indifferent between all prices in the support of the distribution of quotes, including the monopoly price $R(\varepsilon)$. This offer is accepted only if no other dealers who received the RFQ hold the asset, which occurs with probability ϕ_0^{N-1}.

Finally, following similar steps reveals that the flow bid-side profits are

$$\Pi_0(\varepsilon^\star \mid \phi_1) \equiv \left(\frac{\lambda_v}{m}\theta_1 + \frac{\lambda_e N}{m}\phi_1^{N-1}\right)\gamma\,(s - m\phi_1)\int_{-\infty}^{\varepsilon^\star}\frac{\varepsilon^\star-\varepsilon}{(r+\gamma+\chi_1)(\gamma+\lambda_1)}\,\mathrm{d}F(\varepsilon).$$

Taking stock, the difference between the value of acquiring the asset or not is given by the function

$$\Delta(\varepsilon^\star \mid \phi_1) \equiv \mathcal{L}(\varepsilon^\star \mid \phi_1) + \varepsilon_D + \Pi_1(\varepsilon^\star \mid \phi_1) - \Pi_0(\varepsilon^\star \mid \phi_1). \tag{11.}$$

Clearly, the dealer finds it optimal to hold the asset if $\Delta(\varepsilon^\star \mid \phi_1) > 0$, to not hold if $\Delta(\varepsilon^\star \mid \phi_1) < 0$, and is indifferent otherwise.

Proposition 2 (Dealer-optimality schedule). The optimality condition of dealers can be restated

$$\phi_1 = \begin{cases} 0 & & & \geq \\ [0,1] & \text{if} & \varepsilon^\star & = \mathcal{E}(\phi_1), \\ 1 & & & \leq \end{cases} \tag{12.}$$

where $\mathcal{E}(\phi_1) \in \mathbb{R}$ is the unique solution to $\Delta(\varepsilon \mid \phi_1) = 0$. The function $\mathcal{E}(\cdot)$ is continuous, and if the measure of dealers (m) is sufficiently small, $\mathcal{E}(\cdot)$ is strictly decreasing and satisfies $\mathcal{E}(\phi_1) \in (\underline{\varepsilon}, \overline{\varepsilon})$ for all $\phi_1 \in [0, 1]$.

In any equilibrium with active trading on both sides of the market, such that $\phi_1 \in (0, 1)$, dealers must be indifferent between owning the asset (and trading on the ask side) and not owning (and trading on the bid side). According to Proposition 2, this indifference condition implicitly defines a relationship between the fraction ϕ_1 of dealers who hold the asset and the convenience yield of the marginal investor, $\varepsilon^\star = \mathcal{E}(\phi_1)$; we call this relationship the "dealer-optimality schedule." Moreover, Proposition 2 establishes that if m is small, the dealer-optimality schedule is decreasing. Intuitively, if more dealers hold the asset, then they compete more fiercely on the ask side of the

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

market, reducing their profit per trade. For dealers to remain indifferent, they need to make up for those smaller profits by trading a higher volume: Thus, the convenience yield of the marginal customer must fall to expand the set of customers who find it optimal to buy from dealers.

An insight from the dealers' optimality condition is that the relationship between transaction costs and the liquidity yield spread is mediated by two new empirical measures: the ask-side profits, Π_1, and the bid-side profits, Π_0. Importantly, these new measures depend not only on the average electronic and voice transaction costs but also on the order flow per dealer. They can be calculated using information on bid spreads, ask spreads, bond turnover, and the size of the dealer sector.

3.6. Characterization and Properties of Equilibrium

Based on the analysis above, we can define an equilibrium as a pair $(\varepsilon^\star, \phi_1)$ satisfying both the market-clearing condition in Equation 9 and the dealer-optimality condition in Equation 12. Diagrammatically, in an equilibrium with an active dealer sector, $\phi_1 \in (0, 1)$, this amounts to finding the intersection of the two schedules $\phi_1 = \Phi_1(\varepsilon^\star)$ and $\varepsilon^\star = \mathcal{E}(\phi_1)$, as in **Figure 1**. Analytically, this reduces to a fixed point equation $\phi_1 = \Phi_1 \circ \mathcal{E}(\phi_1)$. Proposition 3 establishes equilibrium existence.

Proposition 3. There exists an equilibrium. Moreover, if m is sufficiently small, the equilibrium is unique and the equilibrium fraction of dealers who hold the asset is interior, $\phi_1 \in (0, 1)$.

When m is sufficiently small, the dealer-optimality schedule is strictly decreasing, a sufficient condition for a unique equilibrium. Moreover, comparative statics representing changes in market structure can be studied with a simple diagram in the $(\varepsilon^\star, \phi_1)$ plane, as in **Figure 1**. For example, consider the effects of reducing ε_D, which captures an increase in dealers' balance sheet cost. A reduction of ε_D leaves the market-clearing condition unchanged but shifts the dealers' optimality condition down, as in **Figure 1**.

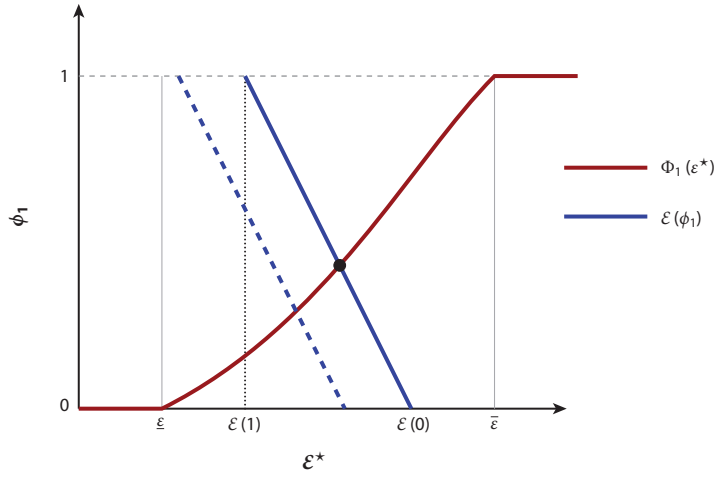


Figure 1

This figure depicts the two equilibrium schedules, and their intersection (*black circle*), when the number of dealers per customer, m, is sufficiently small: the upward sloping market-clearing schedule (*red line*), and the downward sloping dealer-optimality schedule (*solid blue line*). The blue dashed line represents the dealer-optimality schedule following a reduction in ε_D.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

Intuitively, reducing ε_D makes dealers less willing to hold the asset. Hence, holding ϕ_1 constant, restoring the dealers' indifference condition requires the marginal customer's convenience yield $\varepsilon^\star$ to fall, which increases dealers' incentive to hold the asset through three channels. First, since we have normalized the interdealer price $P = 1$, the liquidity yield spread that a dealer earns from holding the asset, ℓ, must rise in equilibrium to offset the decline in $\varepsilon^\star$. Second, a lower value of $\varepsilon^\star$ implies more customers find it optimal to buy the asset, which drives up ask-side profits. Finally, a fall in $\varepsilon^\star$ reduces the set of customers who sell the asset, and thus decreases bid-side profits. When more customers buy from dealers and fewer customers sell, the steady-state equilibrium value of ϕ_1 must fall, as in **Figure 1**. Hence, our model predicts that an increase in balance sheet cost has the natural effect of reducing dealers' total inventory holdings.

Next, we study the impact of increasing balance sheet costs on bid-ask spreads—an important measure of market quality for customers—and on the liquidity yield spread, which determines the impact of liquidity frictions on the cost of capital. To start, we consider trading costs and study the proportional ask spread paid by a customer of type $\varepsilon \geq \varepsilon^\star$ in voice trades,

$$\theta_0 \frac{R(\varepsilon) - P}{P} = \theta_0 \frac{\varepsilon - \varepsilon^\star}{r + \gamma + \chi_0},$$

using Equation 4 and the interdealer market price $P = 1$. One sees that a fall in ε_D has two effects on this spread going in the same direction. First, since $\varepsilon^\star$ decreases, the surplus convenience yield relative to the marginal customer, $\varepsilon - \varepsilon^\star$, increases. Second, since ϕ_1 decreases, competition for customer buyers softens and χ_0 falls. Both effects imply that the ask spread widens, and dealers' profits on the ask side rise. For trades initiated through an electronic RFQ, the effect is stronger because, as noted by Glebkin, Yueshen & Shen (2023), $\Theta(\phi_0)$ rises: Since fewer dealers hold the asset, each dealer is less likely to respond to a request to sell and, thus, respond with less competitive quotes. On the bid side of the market, all the effects discussed above go in the opposite direction: When ε_D falls, it becomes easier to contact willing dealers, dealers make less profits, bid spreads decline for each ε, and the ratio of electronic to voice bid spreads goes down. All in all, this means that the effect on the total bid-ask spread is ambiguous—this finding speaks to the empirical literature that has struggled to identify a meaningful impact of post-GFC regulations on transaction costs.[11]

Finally, turning to the liquidity yield spread, Equation 10 shows that the impact of an increase in balance sheet cost is ambiguous. For example, the decrease in the convenience yield of the marginal investor tends to increase the first component of the liquidity yield spread, but the decrease in the fraction of dealers who hold the asset has the opposite effect, since it decreases the bargaining-adjusted trading intensity on the ask side and increases it on the bid side. To sign the comparative statics and provide quantitative estimates, we now turn to a numerical calibration of our model.

4. QUANTITATIVE ANALYSIS

In this section, we calibrate the model parameters to match key moments in the data. Then, we examine the quantitative impact of changing these parameters in ways that are intended to capture recent changes in the US corporate bond market. Finally, we offer an assessment of the total impact of all three changes occurring concurrently.

[11] In addition, because customer types are unlikely to be observable by an econometrician, the model counterpart of the bid-ask spread measured in the data is the average bid-ask spread across customers. But, since $\varepsilon^\star$ changes with ε_D, the sets of customers who find it optimal to buy, or to sell, change as well, which creates an additional difficulty for this comparative static.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

Table 1 Calibration targets and calibrated parameters[a]

Variable	Value	Empirical target
Discount rate, r	0.05/365	Set to 5% per annum
Voice trade arrival rate, λ_v[b]	0.3096	Estimated voice trade arrival rate[b]
Electronic trade arrival rate, λ_e[b]	0.1512	Estimated electronic trade arrival rate[b]
Number of dealers per customer, m	20/1814	Estimated number of core dealers/unique customers
Bond supply per customer, s	0.3365	Share of corporate bond supply that dealers hold (1.5%)
Preference shock arrival rate, γ	0.00098	Annual bond turnover rate (16.16%)
Number of queried dealers (electronic), N	6	Probability customer seller receives ≥1 reply (99%)
Dealers' bargaining power (voice), θ	0.1232	Average bid-ask spread in voice trades (9.33 bps)
Dealers' convenience yield, ε_D	-2.18×10^{-6}	Annualized liquidity yield spread (28 bps)
Customer convenience yield distribution: $\varepsilon = \mu_0 + \mu_1 X$, where $X \sim$ Beta(α, β)		
Shape parameter, α	5.846	Average ask spread in electronic trades (3.62 bps)
Shape parameter, β	1.9315	Average bid spread in electronic trades (3.89 bps)
Shift parameter, μ_0	−0.009	Probability a customer buyer receives ≥1 reply (97%)
Scaling parameter, μ_1	0.0107	Normalization to make interdealer price $P = 1$

[a]This table presents the baseline numerical values for our model parameters, together with the sources and empirical targets that inform them. Unit of time is 1 day.
[b]Data from Kargar et al. (2025).
Abbreviation: bps, basis points.

4.1. Calibration

We calibrate the model to the US corporate bond market using transaction data from both voice and electronic trades, available from the Trade Reporting and Compliance Engine (TRACE) and MarketAxess, respectively. The period we consider covers January 3, 2017, to March 31, 2021, though we exclude the COVID-19 crisis month of March 2020. Given our assumption of a representative customer and of a representative bond, we restrict attention to a specific market—namely, we consider the sample of institutional-sized trades, exceeding $1 million, for IG bonds. **Table 1** summarizes our choice of baseline parameter values and the associated empirical targets.

The unit of time in the model is set to 1 day, so we set r to deliver an annualized risk-free rate r of 5%. We calibrate the demographics of the market to reflect the structure of the US dealer sector. The measure of dealers is set to $m = 20/1814$, a fraction reflecting the ratio of core corporate bond dealers (Di Maggio, Kermani & Song 2017) to the number of unique customers observed in the MarketAxess data.[12]

The remaining parameters that we need to calibrate are $\{\lambda_e, \lambda_v, s, N, \theta, \gamma, \varepsilon_D\}$, along with the distribution $F(\varepsilon)$; note that, for simplicity, we assume that dealers have the same bargaining power in all voice trades, $\theta_0 = \theta_1 = \theta$. While many of the model-implied moments that we target are jointly determined by several parameters, certain moments are particularly informative about the value of specific parameters. We offer a brief, but intuitive, discussion below.

We set the arrival rates for voice and electronic trades to $\lambda_v = 0.3096$ and $\lambda_e = 0.1512$ to match the empirical execution times estimated by Kargar et al. (2025). Given the (endogenously

[12]To be precise, while Di Maggio, Kermani & Song (2017) define the core group as the top 30 dealers in their baseline specification, they find that their results are robust to narrowing this definition to the top 20 dealers or broadening it to the top 40. We adopt the lower bound of their analysis (20 dealers) to calibrate the size of the dealer sector, since our estimate of the number of customers is also a lower bound.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

determined) probabilities of trading, these arrival rates imply that it takes slightly more than 3 days to execute a trade and that approximately two-thirds of trades occur through voice.

To match the empirical probabilities that a customer seller and customer buyer receive at least one reply to an electronic RFQ—which Kargar et al. (2025) report to be approximately $1 - \phi_1^N \approx 99\%$ for sellers and $1 - (1 - \phi_1)^N \approx 97\%$ for buyers—the number of dealers that receive an RFQ must be $N = 6$ and the (endogenously determined) fraction of dealers holding an asset must be $\phi_1 = 0.4578$. Given ϕ_1 and m, we choose the total supply of assets, $s = 0.3365$, to match the empirical observation that dealers hold 1.5% of the outstanding corporate bond inventory, i.e., to satisfy $0.015 = m\phi_1/s$.

Given ϕ_1, we can also calculate the share of the surplus extracted by dealers in electronic trades, $\Theta(\phi_1)$ and $\Theta(\phi_0)$. We calibrate dealers' bargaining power in voice trades to match the ratio of average electronic bid-ask spreads (Kargar et al. 2025) to average voice bid-ask spreads found in the TRACE data (Cohen et al. 2024). Specifically, the ratio of average ask (bid) spreads in electronic trades to average ask (bid) spreads in voice trades is 0.80 (0.81). Since these ratios are equal to $\Theta(\phi_0)/\theta$ and $\Theta(\phi_1)/\theta$ in our model, they imply different values of $\theta \in \{0.1650, 0.0815\}$. We pick the midpoint of these two numbers, $\theta = 0.1232$.

Turning to the preference parameters, we set the frequency of taste shocks to $\gamma = 0.00098$ to target an annual turnover rate of 16.16% from TRACE, which implies that customers seek to rebalance their positions approximately once every 4 years. The remaining parameters determine the preferences of dealers, ε_D, and of customers, $F(\varepsilon)$, which we take to be a scaled beta distribution—that is, we assume $\varepsilon = \mu_0 + \mu_1 X$, where $X \sim \text{Beta}(\alpha, \beta)$. With this parametric assumption, we have five parameters to determine. We use three target moments and two equilibrium conditions. The three target moments are the liquidity yield spread [28 bps, consistent with the structural decomposition of credit spreads in He & Milbradt (2014)] and the average electronic ask and bid spreads from Kargar et al. (2025) (3.62 and 3.89 bps, respectively). The two equilibrium conditions are that the bond trades at par $R(\varepsilon^\star) = P = 1$ and that the optimality condition $\varepsilon^\star = \mathcal{E}(\phi_1)$, since $\phi_1 \in (0, 1)$. The parameters satisfying these conditions are $\varepsilon_D \approx -8$ bps, $\alpha = 5.846$, $\beta = 1.9315$, $\mu_0 = -3.2807$, and $\mu_1 = 3.9231$, in annualized terms.

4.2. Balance Sheet Costs

We first evaluate the effect of an increase in balance sheet cost from pre- to postcrisis levels. The postcrisis level of convenience yield is that of our calibration, $\varepsilon_D \approx -8$ bps. Fleckenstein & Longstaff (2020) estimate that postcrisis balance sheet costs are 32 bps larger than precrisis. Hence, we take the precrisis level of convenience yield to be 32 bps larger than postcrisis, $\varepsilon_D \approx +24$ bps.

Figure 2*a* reveals that an increase in balance sheet costs increases the liquidity yield spread by approximately 5 bps, which reflects our theoretical observation that the marginal investor's convenience yield, $\varepsilon^\star$, falls. **Figure 2*b*** shows that the impact on the bid-ask spread is muted, which is consistent with empirical studies that have struggled to detect an economically significant impact of regulation on transaction costs (e.g., Bessembinder et al. 2018, Trebbi & Xiao 2019). Intuitively, when ε_D falls, dealers have less incentive to acquire the asset and, in equilibrium, the fraction ϕ_1 of dealers who hold the asset falls. As a result, there is more competition on the bid side of the market and less competition on the ask side: Both average bid and ask prices rise, with little net effect on the bid-ask spread. Interestingly, our results suggest that the effect of balance sheet costs on average bid and ask prices is more pronounced for electronic trades than for voice trades, since a change in ϕ_1 directly affects the number of dealers who respond to RFQs.

Figure 2*a* also reveals a relatively large reduction in bid-side profits alongside a more moderate increase in ask-side profits. One might have guessed that significant changes in ask and bid profits

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

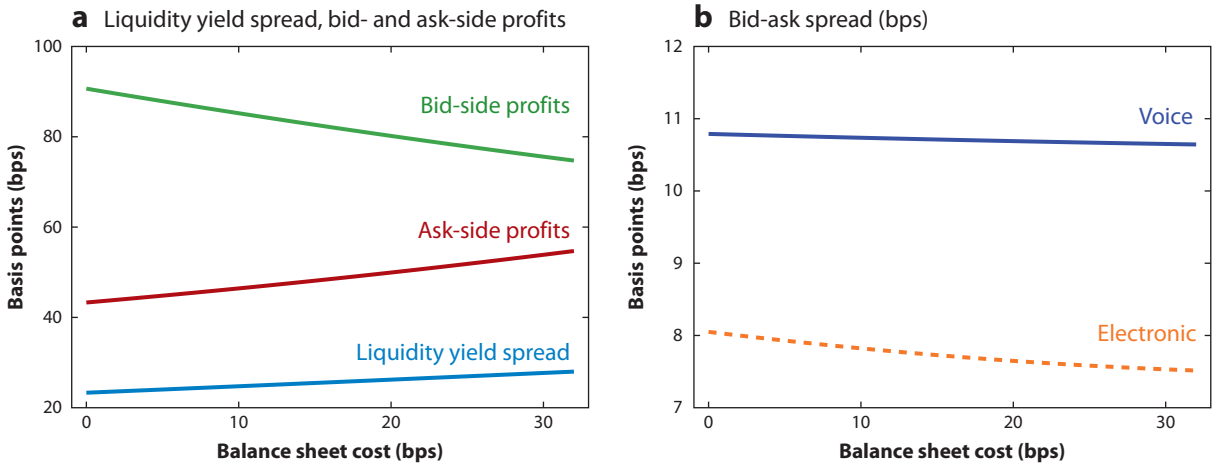


Figure 2

The effects of an increase in balance sheet costs, $-\varepsilon_D$. Panel *a* plots the liquidity yield spread and dealers' ask-side and bid-side profits as functions of the balance sheet cost. As the balance sheet cost rises, the liquidity yield spread and ask-side profits increase, whereas bid-side profits decline. Panel *b* plots the average bid-ask spreads in voice and electronic trades; both decline only modestly, with the decline more pronounced for electronic trades.

would go hand in hand with significant changes in the ask and the bid spreads. However, this is not the case in our quantitative experiment. Indeed, profits are the product of the bid and ask spreads and the corresponding order flow per dealer. In our calibration, the order flow per dealer on either the bid side or the ask side is approximately 14 trades per year, implying that changes in the one-way spread are 14 times smaller than the changes in profits.[13]

4.3. The Rise of Electronic Trading

We now consider a numerical comparative static with respect to the speed of electronic trade, λ_e, generating an increase in the share of electronic trading volume from approximately 20% to 50%, similar to the change in the actual share over the last 10 years (McPartland 2023, 2025). **Figure 3** shows that the increase in λ_e decreases all spreads—the liquidity yield spread, the ask spread, and the bid spread. The intuition is simply that this has the direct effect of reducing search frictions and increasing competition, which brings the OTC market closer to its frictionless counterpart.

Perhaps the most striking observation from **Figure 3** is that the quantitative effect of increasing λ_e is significantly larger than that of rising balance sheet costs: The liquidity yield spread falls by 10 bps and the bid-ask spread by 6 bps. Moreover, one can see that the voice bid-ask spread falls at the same time as the electronic spread; i.e., there are spillovers from electronic to voice trading, as documented by O'Hara & Zhou (2021). Interestingly, we find that the effects are more pronounced on the bid side than on the ask side. Again, the intuition derives from the optimal behavior of dealers: When the fraction of dealers who hold the asset increases, there is less competition on the bid side, which limits the decline in the electronic bid spread by increasing $\Theta(\phi_1)$.

[13] An order flow of 14 trades per year, per dealer, per bond is broadly consistent with other estimates in the literature. For example, Hugonnier, Lester & Weill (2025) report that, on average, a given corporate bond is bought and sold approximately 200 times a year. With 20 core dealers, this translates to roughly 10 trades per bond and per dealer.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

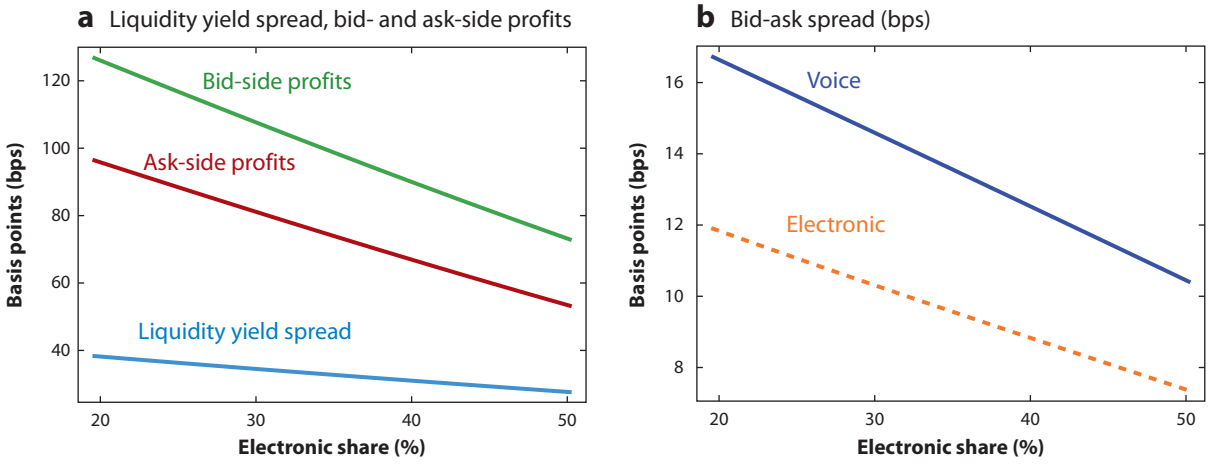


Figure 3

The effect of an increase in share of electronic trading volume, induced by an increase in λ_e. Panel *a* plots the liquidity yield spread and dealers' ask-side and bid-side profits as functions of the electronic trading share; all three decline as the electronic trading share rises. Panel *b* plots the average bid-ask spreads in voice and electronic trades; both decline substantially, indicating spillovers from electronic to voice trading.

4.4. The Rise of Bond Mutual Funds and ETFs

As noted above, the share of corporate bonds held by ETFs and mutual funds has increased from less than 10% to over 20% since 2000. Li & Yu (2026) interpret this as an increase in the importance of short-term investors in the corporate bond market. Indeed, mutual funds and ETFs rebalance frequently in response to benchmark flows, index reconstitutions, and creation/redemption activity—behavior that corresponds in theory to higher Poisson taste-shock arrival rates.

Figure 4 illustrates how increasing the intensity of preference shocks, γ, affects equilibrium outcomes. A higher γ increases turnover, because more investors transition into states in which they want to trade. As turnover rises, the model predicts a systematic increase in the liquidity yield spread: More frequent preference shocks increase the compensation customers require to hold the asset, as in Amihud & Mendelson (1986). From the dealers' optimality condition, the increase in the liquidity yield spread plus the increase in ask-side profits must equal the increase in bid-side profits. Consequently, bid-side profits rise disproportionately relative to ask-side profits. But the increase in bid-side and ask-side profit is, for the most part, due to the increase in order flow per dealer. Hence, bid-ask spreads, both in voice trading and in electronic RFQs, remain essentially flat as γ increases.

4.5. The Combined Effect of the Three Changes in the Market Structure

As a final numerical exercise, we consider the total impact of all three changes in **Table 2**: the increase in balance sheet costs, the rise in the fraction of electronic trading, and the increase in the intensity of preference shocks, γ. We find that the liquidity yield spread increases by less than 3% (less than 1 bps)—an increase that is smaller than the simple sum of the changes calculated above. This calculation reveals that while electronic trading did reduce the cost of capital in the corporate bond market, these gains were almost completely offset by the effect of balance sheet costs and of changes in the demand for liquidity. Bid-ask spreads, however, decreased significantly:

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

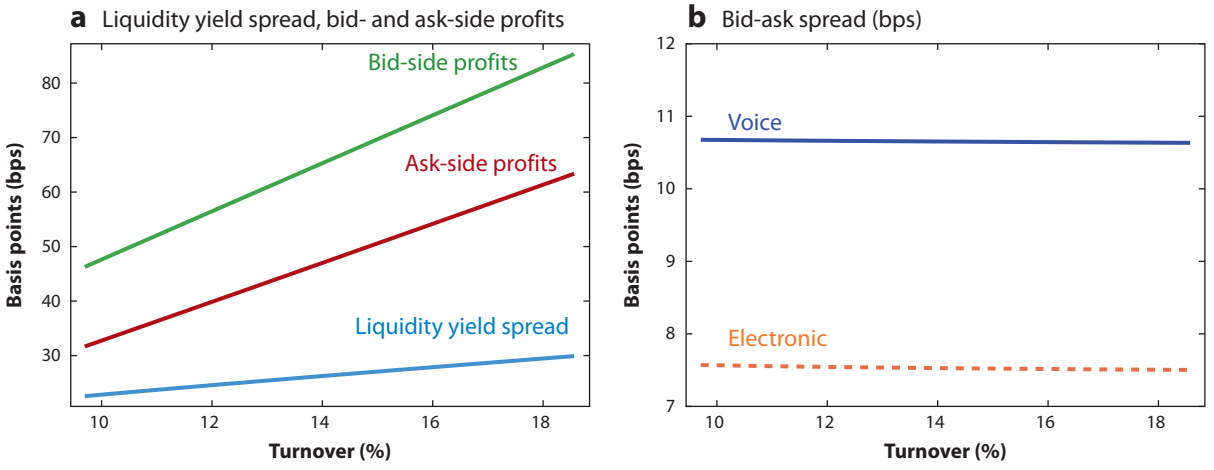


Figure 4

The effects of an increase in turnover induced by an increase in preference shock intensity, γ. Panel *a* plots the liquidity yield spread and dealers' ask-side and bid-side profits as functions of turnover; all three increase, with bid-side profits rising more sharply than ask-side profits. Panel *b* plots the average bid-ask spreads in voice and electronic trades; both remain essentially flat as turnover increases.

We find a decline in average bid-ask spreads of approximately 38% (6.29 bps) for voice trades and 37% (4.41 bps) for electronic trades.

5. CONCLUDING THOUGHTS AND TOPICS FOR FUTURE RESEARCH

A central implication of our analysis is that changes in market structure can significantly affect liquidity by altering dealers' incentives to allocate limited balance sheet space. By studying regulatory balance sheet costs, trading technologies, and the composition of investors within a unified framework, we show that electronic trading substantially lowered the liquidity yield spread but that this reduction in the cost of capital was largely offset by tighter balance sheet constraints and stronger

Table 2 Combined effect analysis: changes in market conditions and spreads[a]

	Change in spread (bps)		
	Liquidity yield	Voice bid-ask	Electronic bid-ask
Panel *a*: Individual effects			
Balance sheet cost (0 → 32 bps)	+5.43	−0.12	+0.68
Electronic share (20% → 50%)	−8.00	−6.20	−4.40
Turnover (11% → 16%)	+6.47	−0.05	−0.00
Panel *b*: Decomposition			
Sum of individual effects	+3.90	−6.37	−3.72
Interaction effect	−3.03	+0.07	−0.68
Combined effect (percent change)	**+ 0.87** (+ 2.8%)	**− 6.29** (− 37.7%)	**− 4.41** (− 37.2%)

[a]This table reports the effects of simultaneous changes in balance sheet costs (+32 bps), electronic share (+30 PP), and turnover (+5 PP) on market spreads. Panel *a* shows the individual effect of each market condition change, holding others constant. Panel *b* decomposes the combined effect into the sum of individual effects and interaction effects.
Abbreviations: bps, basis points; PP, percentage point.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

demand for liquidity. At the same time, transaction costs responded primarily to the change in trading technology, explaining why the bid-ask spread can decline even when the liquidity yield spread remains essentially unchanged.

These results are derived from a framework that is deliberately parsimonious: Its purpose is to provide a unified, tractable analysis of the three major changes in market structure that we study. This parsimony comes at a cost, as several economic forces that are plausibly important in the corporate bond market lie outside the scope of our model. In the remainder of this section, we make these abstractions explicit, discuss what each one delivers and what it leaves out, and connect each to a strand of the literature that confronts it directly. We then close with several broader open questions. Together, these discussions point to potential avenues for future research on dealer-intermediated markets.

5.1. Intensive and Extensive Margins of Intermediation

In our model, dealers can hold either zero or one unit of the asset. This restriction implies that the balance sheet friction operates on the extensive margin: A dealer who already owns the asset can only sell, while a dealer with available balance sheet space can only buy, such that a dealer's current position determines the side of the market on which they can intermediate. This assumption captures extensive margin decisions that dealers make in practice, like whether they should hold a particular bond issue in their inventories. A key benefit of this assumption is that market liquidity is determined, in large part, by one object: the equilibrium fraction ϕ_1 of dealers who hold the asset. This places our analysis within a small set of extensive margin models of multilateral trading, such as Glebkin, Yueshen & Shen (2023), in which dealers' positions are binary and the fraction of dealers who hold the asset, ϕ_1, is the key state variable.

The cost of this assumption, of course, is that the model is silent on dealers' intensive margin decisions: how much bond inventory to carry. Thus, the model can speak neither to how dealers manage the size of their inventories nor to the empirical relationship between trade size and transaction costs, which is a salient feature of the data (Edwards, Harris & Piwowar 2007; Pinter, Wang & Zou 2024). Two intensive margin models illustrate what relaxing this restriction delivers, and what one gives up in return. Cohen et al. (2024), from whom we borrow the assumption that balance sheet space is a key input for intermediation, allow for arbitrary, continuous asset holdings for both dealers and customers. Their model generates richer inventory dynamics, but it no longer admits closed-form solutions and features bilateral rather than multilateral meetings. Maciocco (2026), by contrast, retains multilateral RFQ competition with perfectly divisible assets and endogenous trade sizes and shows that competition among dealers can generate a trade-size discount. To preserve tractability, however, the model abstracts from the balance sheet and inventory channel that are central to our analysis. Developing a framework that combines an intensive inventory margin with multilateral, competitive quoting—and thereby links dealers' balance sheet positions to the size distribution of trades—remains an open challenge that we view as particularly promising.

5.2. Endogenous Dealer Networks

A second abstraction concerns the number of dealers that a customer contacts in an electronic RFQ. Recall that in our model, a customer who trades electronically sends an RFQ to a fixed number of dealers, N, and the random number of quotes they receive is determined by how many of the N contacted dealers are positioned to take the other side of the trade. In our framework, N is a structural parameter, best understood as the breadth of the dealer network that a customer

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

can reach through the platform rather than as a choice made by customers in each RFQ.[14] It is also the parameter that determines the level of competition in each RFQ. A larger N makes a responding dealer more likely to face a competing quote, pushing the winning quote toward the competitive interdealer price, while a smaller N allows dealers to shade their quotes toward the fully discriminating monopoly price, at which the dealer captures the entire surplus from the trade. Because the intensity of competition depends not only on N but also on the endogenous fraction ϕ_1 of dealers who hold the asset, N is simultaneously central to the mechanics of the model and among the harder parameters to discipline directly. Indeed, our calibration identifies N indirectly, from the probability that an RFQ receives at least one reply, rather than from direct observation.

This interpretation also clarifies what is held fixed in Section 4.3: We model the rise of electronic trading as an increase in the speed λ_e at which customers contact electronic platforms, holding N constant. An alternative would be to attribute part of the growth of electronic trading to an increase in N itself, on the grounds that platforms widen the set of dealers a customer can reach. But, in our calibration, this channel alone cannot account for the rise in electronic trading. Indeed, an increase in N makes electronic trading grow only because it increases the probability that a customer receives at least one quote. But this probability is quite large in our calibration, as in the data, so it is rather insensitive to a change in N. While one could of course increase λ_e and N at the same time, we view holding N fixed as a reasonable benchmark. The set of large, active dealers in the corporate bond market has been fairly stable over the period we study, so the first-order change brought about by electronic platforms is plausibly the frequency and ease of multilateral contact rather than an expansion of the set of dealers competing for each trade.

More fundamentally, N is exogenous in our framework, whereas a recent literature treats the effective number of competing dealers as an equilibrium outcome. Wang (2023) shows that when dealers can strategically decline to respond, contacting additional dealers can suppress rather than intensify competition: Each contacted dealer responds with a lower probability and offers a worse price. As a result, a customer optimally contacts only a small number of dealers. Yueshen & Zou (2026) develop a complementary "less can be more" logic, in which a customer invites only a few dealers in order to elicit the greatest total supply of liquidity. In their setting, the optimal number of dealers is pinned down by a competition elasticity that reflects the convexity of dealers' cost of responding. Embedding an endogenous N in our environment would tie the breadth of competition to deeper primitives, namely dealers' response costs and the gains from trade. It would also allow the model to speak directly to whether the growth of electronic trading reflects more frequent contact, a larger set of queried dealers, or a change in how customers strategically choose that set. We leave this extension for future work.

5.3. Customer Heterogeneity

A third abstraction concerns the customer side of the market. In our model, customers are heterogeneous only ex post: They draw idiosyncratic convenience yields that evolve over time, which generates their need to trade. They are otherwise ex ante identical, sharing the same distribution of convenience yields, the same arrival rate of preference shocks, and the same capacity to hold at most one unit of the asset. While this symmetry is convenient and transparent, it requires focusing the calibration on a segment of the market in which customers and bonds are arguably relatively

[14] In practice, electronic RFQ platforms often restrict the set of "disclosed" dealers who receive a customer's inquiry to those dealers with whom the customer had an existing relationship when they joined the platform. Hence, the number of dealers contacted in an RFQ is, to a first approximation, a feature of the customer's existing relationships.

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

homogeneous—namely, institutional-sized trades of IG bonds. Hence, our model cannot speak to the impact of customer heterogeneity on outcomes, such as differences in the size and transaction costs of institutional and retail customers. Introducing ex ante heterogeneity would relax this limitation along several dimensions: Customers could differ in the location or dispersion of their convenience yield distributions, in how frequently they receive preference shocks, and in the quantity of the asset they hold and trade. Such heterogeneity would enrich the model's positive predictions, since different customer types could face systematically different spreads and participation rates, and it would bring the model closer to transaction-level data, in which customers differ widely in size. Models that incorporate this kind of ex ante heterogeneity, such as those of Afonso, Armenter & Lester (2019), Lagos & Navarro (2019), and Pinter, Üslü & Wijnandts (2026), suggest that the cross-sectional distribution of customer characteristics has first-order implications for measured liquidity. Incorporating such heterogeneity into our framework, and revisiting our quantitative analysis with size-weighted moments, is a promising direction for future research.

5.4. Additional Open Questions

Beyond the modeling choices discussed above, our analysis leaves open several broader questions. The first concerns our assumption of a fixed measure of dealers. In the long run, sustained increases in balance sheet costs or persistent reductions in intermediation profits could induce dealer exit, leading to declines in aggregate liquidity that are not captured by our analysis. In addition, our focus on normal times, when buying and selling pressures are roughly balanced, abstracts from crisis episodes, in which customer selling pressure may surge and dealers are expected to absorb inventory in order to stabilize prices. Allowing for endogenous dealer participation and explicitly modeling imbalanced order flow in stress periods are natural extensions that could provide substantial insights into these questions. Fleskes (2025) takes a step in this direction, studying how dealers facing a sudden surge in customer selling pressure choose their intermediation capacity: how quickly to execute orders and when to begin absorbing inventory.

A second set of questions concerns risk and information. We abstract from risk, asymmetric information, and adverse selection, treating the yield spread as arising purely from trading frictions and balance sheet constraints. By contrast, O'Hara & Zhou (2025) highlight the growing role of transparency, electronic price dissemination, and information asymmetries across trading venues in shaping liquidity outcomes, especially during stress episodes. Incorporating risk and informational frictions could offer new insights into the equilibrium relationship between bid-ask spreads, yield spreads, and trading volume, and could even help explain why liquidity provision through electronic platforms appears particularly fragile in periods of heightened uncertainty.

More generally, our results suggest that understanding liquidity in corporate bond markets requires jointly accounting for intermediation capacity, trading technologies, investor composition, and information structure. Developing theoretical frameworks that integrate these dimensions, while preserving tractability and clear empirical counterparts, remains a key challenge for future research.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank the Co-Editors and an anonymous reviewer for their helpful comments and suggestions. The views expressed here are those of the authors and do not necessarily reflect the views of

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

the Federal Reserve Bank of Philadelphia or the Federal Reserve System. All errors are our own responsibility.

LITERATURE CITED

Adrian T, Boyarchenko N, Shachar O. 2017. Dealer balance sheets and bond liquidity provision. *J. Monet. Econ.* 89:92–109

Afonso G, Armenter R, Lester B. 2019. A model of the federal funds market: yesterday, today, and tomorrow. *Rev. Econ. Dyn.* 33:177–204

Amihud Y, Mendelson H. 1986. Asset pricing and the bid-ask spread. *J. Financ. Econ.* 17:223–49

Bao J, O'Hara M, Zhou XA. 2018. The Volcker rule and market making in times of stress. *J. Financ. Econ.* 130:95–113

Bessembinder H, Jacobsen S, Maxwell W, Venkataraman K. 2018. Capital commitment and illiquidity in corporate bonds. *J. Finance* 73:1615–61

Bretscher L, Schmid L, Sen I, Sharma V. 2026. Institutional corporate bond pricing. *Rev. Financ. Stud.* 39(3):605–60

Burdett K, Judd KL. 1983. Equilibrium price dispersion. *Econometrica* 51(4):955–69

Choi J, Huh Y, Shin SS. 2024. Customer liquidity provision: implications for corporate bond transaction costs. *Manag. Sci.* 70(1):187–206

Coalition Greenwich. 2015. *Electronic trading platforms capture more volume in U.S. corporate bond trading*. Press Release, Dec. 8 **https://www.greenwich.com/press-release/electronic-trading-platforms-capture-more-volume-us-corporate-bond-trading**

Cohen A, Kargar M, Lester B, Weill P-O. 2024. Inventory, market making, and liquidity in OTC markets. *J. Econ. Theory* 222:105917

Dannhauser CD, Hoseinzade S. 2022. The unintended consequences of corporate bond ETFs: evidence from the Taper Tantrum. *Rev. Financ. Stud.* 35(1):51–90

Di Maggio M, Kermani A, Song Z. 2017. The value of trading relations in turbulent times. *J. Financ. Econ.* 124(2):266–84

Dick-Nielsen J, Rossi M. 2019. The cost of immediacy for corporate bonds. *Rev. Financ. Stud.* 32:1–41

Duffie D. 2012. Market making under the proposed Volcker rule. Work. Pap. 3118, Graduate School of Business, Stanford University

Duffie D. 2017. *Post-crisis bank regulations and financial market liquidity*. Thirteenth Paolo Baffi Lecture on Money and Finance, Banca d'Italia, Sep. 15

Duffie D, Gârleanu N, Pedersen LH. 2005. Over-the-counter markets. *Econometrica* 73:1815–47

Edwards AK, Harris LE, Piwowar MS. 2007. Corporate bond market transaction costs and transparency. *J. Finance* 62(3):1421–51

Federal Reserve System. 2025. *Table L.213 Corporate and Foreign Bonds*. Flow of Funds, Financial Accounts of the United States - Z.1, Federal Reserve System, updated Sep. 11. **https://www.federalreserve.gov/releases/z1/20250911/html/l213.htm**

Fleckenstein M, Longstaff FA. 2020. Renting balance sheet space: intermediary balance sheet rental costs and the valuation of derivatives. *Rev. Financ. Stud.* 33(11):5051–91

Fleskes N. 2025. Liquidity crises and endogenous dealer capacity in over-the-counter markets. Work. Pap., Rice University

Glebkin S, Yueshen BZ, Shen J. 2023. Simultaneous multilateral search. *Rev. Financ. Stud.* 36:571–614

Goldstein I, Jiang H, Ng DT. 2017. Investor flows and fragility in corporate bond funds. *J. Financ. Econ.* 126(3):592–613

He Z, Milbradt K. 2014. Endogenous liquidity and defaultable bonds. *Econometrica* 82:1443–508

Hendershott T, Madhavan A. 2015. Click or call? Auction versus search in the over-the-counter market. *J. Finance* 70:419–47

Hugonnier J, Lester B, Weill P-O. 2022. Heterogeneity in decentralized asset markets. *Theor. Econ.* 17(3):1313–56

Hugonnier J, Lester B, Weill P-O. 2025. *The Economics of Over-the-Counter Markets: A Toolkit for the Analysis of Decentralized Exchanges*. Princeton University Press

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43

Kargar M, Lester B, Lindsay D, Liu S, Weill P-O, Zúñiga D. 2021. Corporate bond liquidity during the COVID-19 crisis. *Rev. Financ. Stud.* 34:5352–401

Kargar M, Lester B, Plante S, Weill P-O. 2025. Sequential search for corporate bonds. *J. Finance*. Forthcoming

Kozlowski J. 2021. Long-term finance and investment with frictional asset markets. *Am. Econ. J. Macroecon.* 13:411–48

Lagos R, Navarro G. 2019. Monetary operating procedures in the fed funds market: theory and policy analysis. Work. Pap., New York University

Li J, Yu H. 2026. Investor composition and the liquidity component in the U.S. corporate bond market. *J. Finance* 81(2):871–922

Lou D, Pinter G, Üslü S, Walker D. 2025. Yield drifts when issuance comes before macro news. *J. Financ. Econ.* 165:103993

Maciocco A. 2026. Request-for-quote protocols. Work. Pap., University of California, Irvine

Manconi A, Massa M, Yasuda A. 2012. The role of institutional investors in propagating the crisis of 2007–2008. *J. Financ. Econ.* 104(3):491–518

McPartland K. 2023. *Understanding fixed-income markets in 2023*. Tech. Rep., Coalition Greenwich. **https://www.greenwich.com/fixed-income/understanding-fixed-income-markets-2023**

McPartland K. 2025. *Corporate bond e-trading looks for next growth catalyst*. Tech. Rep., Coalition Greenwich. **https://www.greenwich.com/market-structure-technology/corporate-bond-e-trading-looks-next-growth-catalyst**

O'Hara M, Zhou XA. 2021. The electronic evolution of corporate bond dealers. *J. Financ. Econ.* 140(2):368–90

O'Hara M, Zhou XA. 2025. US corporate bond markets: bigger and (maybe) better? *J. Econ. Perspect.* 39(2):215–34

Pinter G, Üslü S, Wijnandts JC. 2026. Comparing search and intermediation frictions across markets. Work. Pap., Johns Hopkins University

Pinter G, Wang C, Zou J. 2024. Size discount and size penalty: trading costs in bond markets. *Rev. Financ. Stud.* 37:2156–90

Thakor AV. 2012. *The economic consequences of the Volcker rule*. Rep., US Chamber of Commerce Center for Capital Market Competitiveness

Timmer Y. 2018. Cyclical investment behavior across financial institutions. *J. Financ. Econ.* 129(2):268–86

Trebbi F, Xiao K. 2019. Regulation and market liquidity. *Manag. Sci.* 65(5):1949–68

Wang C. 2023. The limits of multi-dealer platforms. *J. Financ. Econ.* 149:434–50

Yueshen BZ, Zou J. 2026. When can less be more? Inducing dealer responses in RFQ platforms. Work. Pap., INSEAD

Downloaded from www.annualreviews.org. Guest (guest) IP: 91.24.127.21 On: Wed, 04 Nov 2026 15:00:43