FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

The Evolution of the Corporate Bond Market
A Theoretical Analysis

Mahyar Kargar
University of Illinois Urbana–Champaign

Benjamin Lester
Federal Reserve Bank of Philadelphia

Semih Üslü
Johns Hopkins University

Pierre-Olivier Weill
UCLA, NBER, and CEPR

WP 26-27

PUBLISHED
May 2026

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2026.27

The Evolution of the Corporate Bond Market: A Theoretical Analysis*

Mahyar Kargar[†] Benjamin Lester[‡] Semih Üslü[§] Pierre-Olivier Weill[¶]

April 28, 2026

Abstract

We develop a model of a dealer-intermediated over-the-counter market designed to study three major changes in the structure of the U.S. 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 ETFs. 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.

Keywords: Over-the-counter markets, corporate bond market liquidity, dealer intermediation, balance sheet costs, electronic trading platforms.

JEL Classification: G11, G12, G14, G24, D83.

*The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. All errors are our own responsibility.

[†]University of Illinois Urbana-Champaign. Email: kargar@illinois.edu

[‡]Federal Reserve Bank of Philadelphia. Email: benjamin.lester@phil.frb.org

[§]Johns Hopkins University. Email: semihuslu@jhu.edu

[¶]UCLA, NBER and CEPR. Email: poweill@econ.ucla.edu

1 Introduction

The corporate bond market, the primary source of long-term borrowing for large American companies, has changed in a number of 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 easy 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) that have starkly different trading needs.[3]

A recent literature has extensively documented 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 OTC market for a representative bond 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.

[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 investment-grade 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 shows that from 2000Q1 to 2025Q2, 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%.

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 inter-dealer 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, which ultimately determines the real cost of capital and, hence, the macroeconomic effects of secondary market frictions.

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, Jacobsen, Maxwell, and Venkataraman, 2018; Trebbi and 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 and Yu (2025). 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 inter-dealer market—incurring exogenous holding costs to service ask-side demand—or they can reserve capacity to accommodate bid-side selling by 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 optimality condition of dealers requires that the sum of ask-side profits and 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—the product of the ask and bid spread with the order flow per dealer. 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 profits, liquidity yield spread, and bid profits. 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.

To start, we consider an increase in balance sheet costs of 32 bps, as estimated by Fleckenstein and 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 about 5 basis points, and the fraction of bonds held in the dealer sector falls by approximately 6 percent. This means that there is less competition among dealers who own assets, so that ask prices rise and ask-side profits increase by about 11 basis points. At the same time, since more dealers have balance sheet space, bid prices become more competitive and bid-side profits fall by about 16 basis points. 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 speed at which customers contact electronic platforms, so that our model generates an increase in the share of electronic trading from about 20 percent to 50 percent, 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 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 towards 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.

A brief note on the literature. Our market setting is a variation of the canonical search-theoretic model of OTC markets developed by Duffie, Gârleanu, and Pedersen (2005); we refer the reader to Hugonnier, Lester, and 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, and Shen (2023), we incorporate multilateral requests for quotes to study the emergence of electronic trading platforms, though we extend their theoretical market setting in two important ways. First, we introduce a competitive inter-dealer market to study the determination of an equilibrium inter-dealer price, which is conceptually similar to the reference price calculated in empirical studies (see, e.g., Choi, Huh, and Shin, 2024). This facilitates the analysis of the liquidity yield spread as well as 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, and 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, Kargar, Lester, and Weill (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 inter-dealer 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 paper proceeds as follows. Section 2 offers a brief description of the three recent changes in the U.S. 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 concludes and discusses potential avenues for future research.

2 Background

In this section, we briefly describe the three major changes in the corporate bond market that we study using the theoretical framework developed below. 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 and Zhou (2025).

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.[4]

Following the implementation of these regulatory reforms, there is evidence that affected dealers curtailed their market-making activities, and non-affected dealers did not fully offset this contraction. For one, the dealer sector as a whole reduced the level of assets on their balance sheets: according to the Flow of Funds, the share of outstanding corporate bonds held by security brokers and dealers fell from 2-3% in 2006 to less than 1% in 2018.[5] 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.[6]

Highlighting the importance of dealers' balance sheet space in providing liquidity, several studies have documented that bond-level liquidity comoves with intermediaries' balance-sheet tight-

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

[5]Source: The Financial Accounts of the United States (Flow of Funds) Table L.213. The data also include foreign bonds, privately issued mortgage-backed securities, and other privately issued asset-backed bonds.

[6]See, for example, Bao, O'Hara, and Zhou (2018), Bessembinder et al. (2018), Trebbi and Xiao (2019), Dick-Nielsen and Rossi (2019), and Choi, Huh, and Shin (2024).

ness: when constraints bind, transaction costs rise and dealers' willingness to take inventory falls, consistent with a state-dependent ‘‘price of immediacy.’’[7] Liquidity conditions during the COVID-19 episode provided further evidence that dealers had become less willing to ‘‘lean against the wind;’’ for example, Kargar, Lester, Lindsay, Liu, Weill, and Zúñiga (2021) document that the cost of risky-principal trades (where dealers absorb bonds onto their balance sheets) spiked in March 2020 and trading migrated to agency trades (where dealers match customer buyers and sellers without holding the asset themselves).

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 at once via requests for quotes (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 investment grade (IG) bonds grew from about 20% in 2015 to about 50% in 2025. As electronic trading gained market share, O'Hara and Zhou (2021) document that transaction costs fell in both electronic and voice trades.

It is important to note that voice-based trading remains a significant fraction of trading volume, particularly in certain segments of the corporate bond market.[8] Also dealers' inventory holdings remain a key determinant of their willingness and ability to trade with customers. In particular, Kargar et al. (2025) 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).

The rise of bond mutual funds and ETFs. A third important development in the corporate bond market since the GFC relates to the composition of investors who own the bonds. In particular,

[7] See, e.g., Adrian, Boyarchenko, and Shachar (2017) and Lou, Pinter, Üslü, and Walker (2025).

[8] In particular, electronic trading is most pronounced for smaller, investment-grade trades, while large blocks and lower-rated issues remain predominantly negotiated by voice; see, e.g., Hendershott and Madhavan (2015), O'Hara and Zhou (2021), and Kargar, Lester, Plante, and Weill (2025).

ownership 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 the Flow of Funds, 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.[9]

This change in the ownership structure of corporate bonds has a variety of potential effects on the demand for liquidity. The literature following Goldstein, Jiang, and Ng (2017) has emphasized that the rise of open-ended funds and ETFs has increased the sensitivity and magnitude of demand shocks (for liquidity), 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 (Timmer, 2018; Bretscher, Schmid, Sen, and Sharma, 2025) than traditional buy-and-hold investors, while ETFs tend to attract investors with greater liquidity demand (Manconi, Massa, and Yasuda, 2012; Dannhauser and Hoseinzade, 2022).

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—both implicit and explicit—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. And 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

[9]See Table L.213 from the Flow of Funds, and Li and Yu (2025) for more details.

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.

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 one 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 $P = 1$.[10] 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 will refer to ℓ as the ''liquidity yield spread'' since, in our framework, the yield spread only derives 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 will 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 semi-centralized: customers must search for dealers in order to trade, while dealers can trade together in a frictionless competitive market. But we depart from the

[10]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.

standard semi-centralized model described in Chapter 4 of Hugonnier, Lester, and Weill (2025) by incorporating model elements from Glebkin, Yueshen, and 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 request for quotes (RFQ), the investor contacts $N > 1$ dealers, each of whom may respond with an offer. Importantly, when a dealer who receives an electronic RFQ responds, they do not know how many other dealers will also respond, as in Burdett and Judd (1983).

Second, we assume that a dealer must choose their inventory holdings, which determines the direction of their trade, should they come into contact with a customer. That is, 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 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 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 meeting customers, their reservation value is equal to the price of the asset in the perfectly competitive inter-dealer market, P.

Voice trades. Consider a bilateral meeting via voice between a dealer who owns one unit of the asset and a customer non-owner 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 non-owner 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

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

Electronic trades. Consider now an electronic multilateral meeting, in which a customer non-owner 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, and 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 and 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 Chapter 8 of Hugonnier, Lester, and Weill (2025). The lowest quoted price is strictly greater than the competitive price, P, but the highest price is equal to the monopoly price, $R(\varepsilon)$. This 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

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

3.3 The Customer's Problem

Consider a customer non-owner with convenience yield ε. The Hamilton-Jacobi-Bellman (HJB) equation for this customer is:

$$\begin{aligned} rV(0,\varepsilon) = {} & \gamma\left(\int V(0,x)dF(x) - V(0,\varepsilon)\right) \\ & + \lambda_v\,(1-\phi_0)\left[V(1,\varepsilon) - V(0,\varepsilon) - \theta_0 R(\varepsilon) - (1-\theta_0)P\right]^+ \\ & + \lambda_e(1-\phi_0^N)\left[V(1,\varepsilon) - V(0,\varepsilon) - \Theta(\phi_0)R(\varepsilon) - (1-\Theta(\phi_0))P\right]^+, \end{aligned}$$

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 $\lambda_v(1-\phi_0)$ 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:

$$\begin{aligned} rV(0,\varepsilon) &= \gamma\left(\int V(0,x)dF(x) - V(0,\varepsilon)\right) \\ &\quad + \lambda_v\,(1-\phi_0)\,(1-\theta_0)\Big[R(\varepsilon)-P\Big]^+ + \lambda_e(1-\phi_0^N)\,(1-\Theta(\phi_0))\,\Big[R(\varepsilon)-P\Big]^+. \end{aligned} \tag{1}$$

This reveals that a customer's value function can be calculated ''as if'' she could trade directly in the inter-dealer 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:

$$\begin{aligned} rV(1,\varepsilon) &= r + \ell + \varepsilon + \gamma\left(\int V(1,x)dF(x) - V(1,\varepsilon)\right) \\ &\quad + \lambda_v\,(1-\phi_1)\,(1-\theta_1)\Big[P-R(\varepsilon)\Big]^+ + \lambda_e(1-\phi_1^N)\,(1-\Theta(\phi_1))\,\Big[P-R(\varepsilon)\Big]^+. \end{aligned} \tag{2}$$

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

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

where

$$\chi_q \equiv \lambda_v\left(1-\phi_q\right)\left(1-\theta_q\right) + \lambda_e(1-\phi_q^N)\left(1-\Theta(\phi_q)\right)$$

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 semi-centralized model, the bargaining-adjusted intensities are different on the ask and the bid sides 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, $q \in \{0, 1\}$ 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, and hence shifts the distribution towards 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 $R(\varepsilon^\star) = P$. It follows from results in Hugonnier, Lester, and Weill (2025) that the reservation value function is strictly increasing in ε, differentiable everywhere except at $\varepsilon^\star$, with derivative:

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

3.4 Distributions and Market-clearing

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 non-owner buys from a dealer if $\varepsilon > \varepsilon^\star$, while a customer owner sells to a dealer if $\varepsilon < \varepsilon^\star$.[11] Hence, in a steady-state equilibrium, the inflow and outflow of customers in state (q, ε) must be equal, so that

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

The left-hand side of (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 non-owners 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 (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 (5), we obtain:

$$\begin{aligned} d\Psi_0(\varepsilon) &= \frac{\lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}} + \gamma(1 - s + m\phi_1)}{\gamma + \lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}}} \, dF(\varepsilon), \\ d\Psi_1(\varepsilon) &= \frac{\lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \gamma \left(s - m\phi_1\right)}{\gamma + \lambda_0 \mathbb{I}_{\{\varepsilon \geq \varepsilon^\star\}} + \lambda_1 \mathbb{I}_{\{\varepsilon < \varepsilon^\star\}}} \, dF(\varepsilon). \end{aligned} \tag{7}$$

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:

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

[11] 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.

Using the explicit expressions for λ_0 and λ_1, we obtain:

$$\frac{\gamma\,(s-m\phi_1)}{\gamma+\lambda_1}F(\varepsilon^\star)+\frac{\gamma\,(s-m\phi_1)+\lambda_0}{\gamma+\lambda_0}\left(1-F(\varepsilon^\star)\right)+m\phi_1=s. \tag{8}$$

The following Proposition establishes that the market-clearing condition (8) 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 (8) 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 $\varepsilon \in (\underline{\varepsilon}, \overline{\varepsilon})$, 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 non-owners 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 it is impossible to buy, 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 or not to acquire a unit of the asset in the inter-dealer 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

$$\Delta \equiv r + \ell + \varepsilon_D - rP + \Pi_1 - \Pi_0 = \ell + \varepsilon_D + \Pi_1 - \Pi_0,$$

where the equality follows from our normalizing assumption that the bond trades at par ($P = 1$) in the inter-dealer 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

$$\mathcal{L}\left(\varepsilon^\star \mid \phi_1\right) \equiv -\varepsilon^\star - \frac{\gamma}{r+\gamma+\chi_0} \int_{\varepsilon^\star}^{+\infty} \left(\varepsilon - \varepsilon^\star\right) \, dF(\varepsilon) + \frac{\gamma}{r+\gamma+\chi_1} \int_{-\infty}^{\varepsilon^\star} \left(\varepsilon^\star - \varepsilon\right) \, dF(\varepsilon). \quad (9)$$

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 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:

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

This equality is established by substituting the explicit formula for $d\Psi_0(\varepsilon)$ from equation (7) and applying equation (4), noting that $R(\varepsilon) - P = \int_{\varepsilon^\star}^{\varepsilon} R'(x)\, dx = (\varepsilon - \varepsilon^\star)/(r+\gamma+\chi_0)$. 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 only accepted 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:

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

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

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

Clearly the dealer finds it optimal to hold the asset if $\Delta(\varepsilon^\star \mid \phi_1) > 0$, not to 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 re-stated*

$$\phi_1 = \begin{cases} 0 & \geq \\ [0,1] \quad \text{if} \quad \varepsilon^\star & = \mathcal{E}(\phi_1), \\ 1 & \leq \end{cases} \tag{10}$$

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, so 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, the Proposition 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 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 not only depend 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 (8) and the dealer optimality condition in equation (10). Diagrammatically, in an equilibrium with an active dealer sector, $\phi_1 \in (0, 1)$, this boils down 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 $\phi_1 = \Phi_1 \circ \mathcal{E}(\phi_1)$. The following Proposition 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

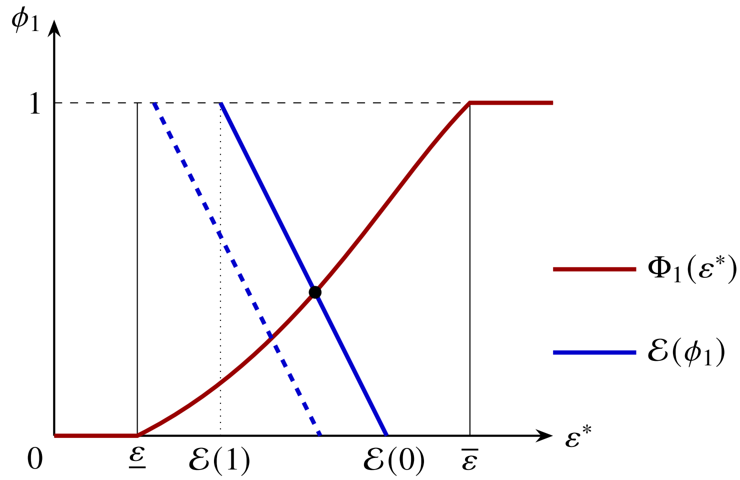

Figure 1. This figure depicts the two equilibrium schedules, and their intersection, when m is sufficiently small: the upward sloping market-clearing schedule, and the downward sloping dealer optimality schedule. The dashed line represents the dealer optimality schedule following a reduction in ε_D.

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 the dashed schedule in Figure 1.

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 inter-dealer 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 the diagram. 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 cost on trading costs, an important measure of market quality for customers, and 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 inter-dealer 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, and 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* of 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.[12]

Finally, turning to the liquidity yield spread, equation (9) 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.

[12] 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.

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 U.S. corporate bond market. Finally, we offer an assessment of the *total* impact of all three changes occurring concurrently.

4.1 Calibration

We calibrate the model to the U.S. corporate bond market using transaction data from both voice and electronic trades, available from TRACE and MarketAxess, respectively. The period we consider covers from January 3, 2017 to March 31, 2021, though we exclude the COVID-19 crisis month of March 2020. We restrict our sample to trades of investment-grade (IG) bonds with a transaction size exceeding \$1 million. Table 1 summarizes our choice of baseline parameter values as well as the associated empirical targets.

The unit of time in the model is set to one 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 U.S. 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, and Song, 2017) to the number of unique customers observed in the MarketAxess data.[13]

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 a number of parameters, there are certain moments that 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 $\lambda_e = 0.1512$ and $\lambda_v = 0.3096$ to

[13]To be precise, while Di Maggio, Kermani, and 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.

Variable	Value	Empirical target
Discount rate, r	$0.05/365$	Set to 5% per annum
Voice trade arrival rate, λ_v	0.3096	Source: Kargar et al. (2025)
Electronic trade arrival rate, λ_e	0.1512	Source: Kargar et al. (2025)
Number of dealers per customer, m	20/1814	Est. no. of core dealers / no. of 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 (elec.), N	6	Prob. customer-seller receives ≥ 1 reply (99%)
Dealers' bargaining power (voice), θ	0.1232	Avg. bid-ask spread in voice trades (9.33 bps)
Dealers' convenience yield, ε_D	-2.18×10^{-6}	Annualized liquidity yield spread (28 bps)
Customers convenience yield distribution		$\varepsilon = \mu_0 + \mu_1 X$, *where* $X \sim \text{Beta}(\alpha, \beta)$
Shape parameter, α	5.846	Avg. ask spread in electronic trades (3.62 bps)
Shape parameter, β	1.9315	Avg. bid spread in electronic trades (3.89 bps)
Shift parameter, μ_0	−0.009	Prob. a customer-buyer receives ≥ 1 reply (97%)
Scaling parameter, μ_1	0.0107	Normalization to make inter-dealer price $P = 1$

Table 1. Calibration targets and calibrated parameters. 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 a day.

match the empirical execution times estimated by Kargar et al. (2025). Given the (endogenously determined) probabilities of trading, these arrival rates imply that it takes slightly more than three 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 four 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 and Milbradt (2014)) and the average electronic ask and bid spreads from Kargar et al. (2025) (3.62 bps and 3.89 bps, respectively). The two equilibrium conditions are that the bond trades at par $R(\varepsilon^\star) = P = 1$, and 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 start with a comparative static with respect to balance sheet costs. Since our calibration matches data from the post-crisis years, we evaluate the effect of a return to pre-crisis regulation by increasing the convenience yield of dealers by about 32 basis points (bps), following estimates of Fleckenstein and Longstaff (2020). In Figure 2, a balance sheet cost of 32 bps refers to our calibration, where $\varepsilon_D \approx -8$ bps per year. Accordingly, a baseline of zero balance sheet costs corresponds to $\varepsilon_D \approx +24$ bps per year, representing the dealer's convenience yield in the pre-crisis regulatory environment.

The left panel of Figure 2 reveals that an increase in balance sheet costs increases the liquidity yield spread, by about 5 bps, which reflects our theoretical observation that the marginal investor's convenience yield, $\varepsilon^\star$, falls. The right panel shows that the impact on the bid-ask spread is muted, which is consistent with empirical studies that have struggled detecting an economically significant

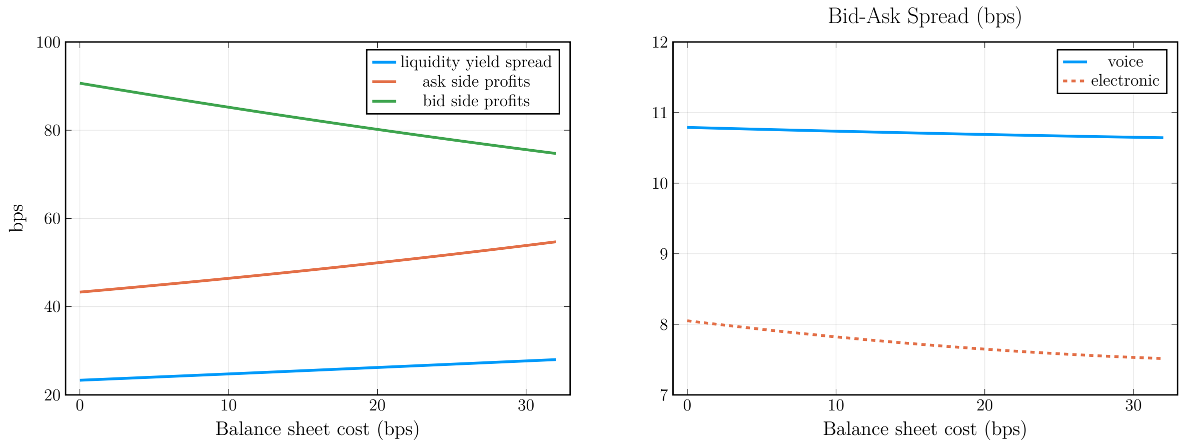

Figure 2. The effects of an increase in balance sheet costs, $-\varepsilon_D$.

impact of regulation on transaction costs (e.g., Bessembinder et al., 2018; Trebbi and 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: average bid and ask prices both 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.

The left panel of Figure 2 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* 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 or the ask-side is of about 14 trades per year, implying that changes in the one-way spread are 14 times smaller than changes in profits.[14]

[14] 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, and 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.

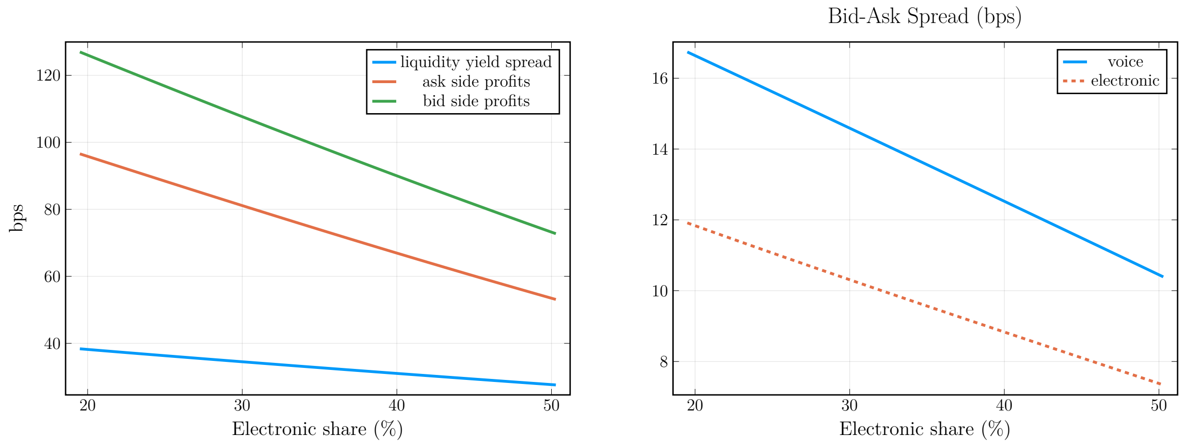

Figure 3. The effect of an increase in share of electronic trading volume, induced by an increase in λ_e.

4.3 The Rise of Electronic Trading

We now consider a numerical comparative static with respect to the speed of electronic trade, λ_e, creating an increase in the electronic share of the trading volume from about 20 percent to 50 percent, 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 the rise in balance sheet cost: 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 and Zhou (2021). Interestingly, we find that the effects are more pronounced on the bid-side, relative to the ask. 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 $\Theta(\phi_1)$.

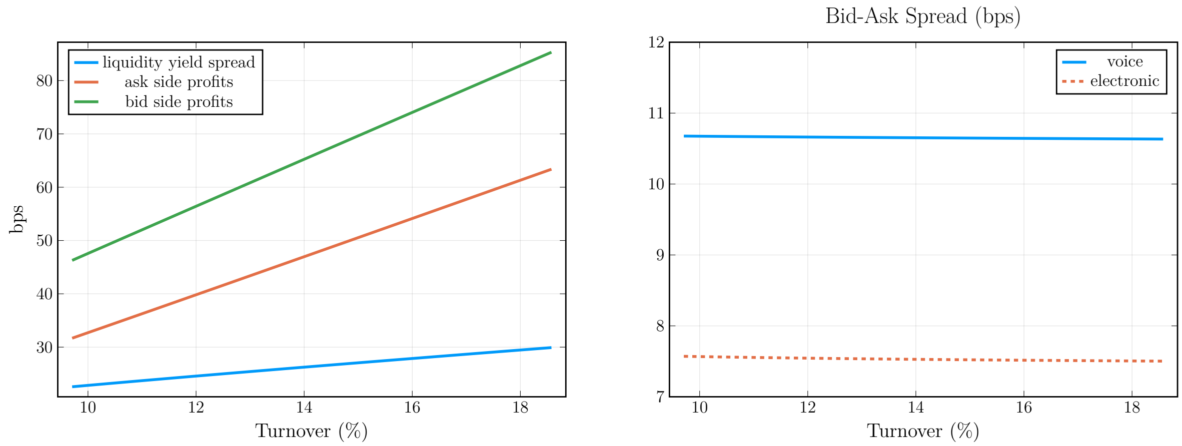

Figure 4. The effects of an increase in turnover induced by an increase in preference shock intensity, γ.

4.4 The Rise of Bond Mutual Funds and ETFs

As noted earlier, the share of corporate bonds held by ETFs and mutual funds has increased from less than 10% to over 20% since 2000. Li and Yu (2025) 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 the theory to higher Poisson taste-shock arrival rates.

Figure 4 illustrates how increasing the preference-shock intensity γ affects equilibrium outcomes in the OTC market. 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 and 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

	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	**+0.87**	**−6.29**	**−4.41**
(Percent Change)	(+2.8%)	(−37.7%)	(−37.2%)

Table 2. Combined Effect Analysis: Changes in Market Conditions and Spreads.
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.

As a final numerical exercise, we consider the total impact of all three changes: the increase in balance sheet costs, the rise in the fraction of electronic trading, and the increase in the frequency of trading. 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: we find a decline in average bid-ask spreads of about 38% (6.29 bps) for voice trades, and 37% (4.41 bps) for electronic trades.

5 Concluding Thoughts and Open Questions

A central implication of our analysis is that changes in market structure affect liquidity primarily through 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 improved allocative efficiency by lowering the liquidity yield spread, but the associated efficiency gains were largely offset by tighter balance sheet constraints and heightened liquidity demand from funds. At the same time, transactional efficiency, as measured by transaction costs, responded primarily to the change in trading technology, which explains why observed bid-ask spreads can decline even as the economic cost of illiquidity remains essentially unchanged.

An important abstraction underlying these results is 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.

A second set of open questions concerns mechanisms that lie outside the scope of our model but are central to the broader evolution of the corporate bond market. We abstract from risk, asymmetric information, and adverse selection, treating the liquidity yield spread as arising purely from trading frictions and balance sheet constraints. By contrast, O'Hara and 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 complement the equilibrium relationship between bid-ask spreads, yield spreads, and trading volume that our model uncovers and may 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 technology, 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.

References

Adrian, Tobias, Nina Boyarchenko, and Or Shachar, 2017, Dealer balance sheets and bond liquidity provision, *Journal of Monetary Economics* 89, 92–109.

Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223–249.

Bao, Jack, Maureen O'Hara, and Xing (Alex) Zhou, 2018, The volcker rule and market making in times of stress, *Journal of Financial Economics* 130, 95–113.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, *Journal of Finance* 73, 1615–1661.

Bretscher, Lorenzo, Lukas Schmid, Ishita Sen, and Varun Sharma, 2025, Institutional corporate bond pricing, *Review of Financial Studies* forthcoming.

Burdett, Kenneth, and Kenneth L. Judd, 1983, Equilibrium price dispersion, *Econometrica* 51, 955–969.

Choi, Jaewon, Yesol Huh, and Sean Seunghun Shin, 2024, Customer liquidity provision: Implications for corporate bond transaction costs, *Management Science* 70, 187–206.

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

Cohen, Assa, Mahyar Kargar, Benjamin Lester, and Pierre-Olivier Weill, 2024, Inventories, market making and liquidity in OTC markets, *Journal of Economic Theory* 222, 105917.

Dannhauser, Caitlin D., and Saeid Hoseinzade, 2022, The unintended consequences of corporate bond etfs: Evidence from the taper tantrum, *Review of Financial Studies* 35, 51–90.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song, 2017, The value of trading relations in turbulent times, *Journal of Financial Economics* 124, 266–284.

Dick-Nielsen, Jens, and Marco Rossi, 2019, The cost of immediacy for corporate bonds, *Review of Financial Studies* 32, 1–41.

Duffie, Darrell, 2012, Market making under the proposed Volcker rule, Technical report, Graduate School of Business, Stanford University.

Duffie, Darrell, 2017, *Post-crisis bank regulations and financial market liquidity* (Banca d'Italia).

Duffie, Darrell, Nicolae Gârleanu, and Lasse H. Pedersen, 2005, Over-the-counter markets, *Econometrica* 73, 1815–1847.

Fleckenstein, Matthias, and Francis A. Longstaff, 2020, Renting balance sheet space: Intermediary balance sheet rental costs and the valuation of derivatives, *Review of Financial Studies* 33, 5051–5091.

Glebkin, Sergei, Bart Zhou Yueshen, and Ji Shen, 2023, Simultaneous multilateral search, *Review of Financial Studies* 36, 571–614.

Goldstein, Itay, Hao Jiang, and David T. Ng, 2017, Investor flows and fragility in corporate bond funds, *Journal of Financial Economics* 126, 592–613.

He, Zhiguo, and Konstantin Milbradt, 2014, Endogenous liquidity and defaultable bonds, *Econometrica* 82, 1443–1508.

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

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, 2022, Heterogeneity in decentralized asset markets, *Theoretical Economics* 17, 1313–1356.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, 2025, *The Economics of Over-the-Counter Markets: A Toolkit for the Analysis of Decentralized Exchanges* (Princeton University Press).

Kargar, Mahyar, Benjamin Lester, David Lindsay, Shuo Liu, Pierre-Olivier Weill, and Diego Zúñiga, 2021, Corporate bond liquidity during the COVID-19 crisis, *Review of Financial Studies* 34, 5352–5401.

Kargar, Mahyar, Benjamin Lester, Sébastien Plante, and Pierre-Olivier Weill, 2025, Sequential search for corporate bonds, *Journal of Finance* forthcoming.

Li, Jian, and Haiyue Yu, 2025, Investor composition and the liquidity component in the U.S. corporate bond market, *Journal of Finance* forthcoming.

Lou, Dong, Gabor Pinter, Semih Üslü, and Danny Walker, 2025, Yield drifts when issuance comes before macro news, *Journal of Financial Economics* 165, 103993.

Manconi, Alberto, Massimo Massa, and Ayako Yasuda, 2012, The role of institutional investors in propagating the crisis of 2007–2008, *Journal of Financial Economics* 104, 491–518.

McPartland, Kevin, 2023, Understanding fixed-income markets in 2023, Technical report, Coalition Greenwich.

McPartland, Kevin, 2025, Corporate bond e-trading looks for next growth catalyst, Technical report, Coalition Greenwich.

O'Hara, Maureen, and Xing (Alex) Zhou, 2021, The electronic evolution of corporate bond dealers, *Journal of Financial Economics* 140.

O'Hara, Maureen, and Xing (Alex) Zhou, 2025, US corporate bond markets: Bigger and (maybe) better?, *Journal of Economic Perspectives* 39, 215–234.

Thakor, Anjan V., 2012, The economic consequences of the Volcker rule, Report by the US Chamber of Commerce Center for Capital Market Competitiveness.

Timmer, Yannick, 2018, Cyclical investment behavior across financial institutions, *Journal of Financial Economics* 129, 268–286.

Trebbi, Francesco, and Kairong Xiao, 2019, Regulation and market liquidity, *Management Science* 65, 1949–1968.

Appendix

A Omitted Proofs

A.1 Proof of Proposition 1

We first establish that the left-hand side of (8) is strictly increasing in ϕ_1. This can be seen by noting that the left-hand side depends on ϕ_1 in two ways: directly and indirectly through its effect on λ_k. We see that the direct effect is positive by taking the partial derivative of (8) with respect to ϕ_1:

$$-\frac{\gamma m}{\gamma+\lambda_1}F(\varepsilon^\star) - \frac{\gamma m}{\gamma+\lambda_0}\left(1-F(\varepsilon^\star)\right) + m > -mF(\varepsilon^\star) - m\left(1-F(\varepsilon^\star)\right) + m = 0.$$

The indirect effects are also positive. Indeed, the left-hand side is decreasing in λ_1, but λ_1 is decreasing in ϕ_1. Likewise, the left-hand side is increasing in λ_0, but λ_0 is decreasing in ϕ_0, hence increasing in ϕ_1.

Now, evaluating the left-hand side of the equation at $\phi_1 = 0$ leads to

$$\frac{\gamma s}{\gamma+\lambda_v+\lambda_e}F(\varepsilon^\star) + s(1-F(\varepsilon^\star)) \leq s,$$

with a strict inequality if $\varepsilon^\star > \underline{\varepsilon}$, and an equality otherwise. Likewise, evaluating at $\phi_1 = 1$:

$$(s-m)\,F(\varepsilon^\star) + \frac{\gamma(s-m)+\lambda_e+\lambda_v}{\gamma+\lambda_e+\lambda_v}\left(1-F(\varepsilon^\star)\right) + m \geq s,$$

with a strict inequality if $\varepsilon^\star < \overline{\varepsilon}$, and an equality otherwise. An application of the Intermediate Value Theorem delivers existence and uniqueness of a solution, as well as the stated bounds. The usual proof of continuity follows because (8) is continuous in $(\varepsilon^\star, \phi_1)$ and the solution is unique. Finally, to establish that $\Phi_1(\cdot)$ is increasing, we take the derivative of (8) with respect to $\varepsilon^\star$ and obtain:

$$\begin{aligned} & dF(\varepsilon^\star)\left(\frac{\gamma(s-m\phi_1)}{\gamma+\lambda_1} - \frac{\gamma(s-m\phi_1)+\lambda_0}{\gamma+\lambda_0}\right) \\ = & -\frac{dF(\varepsilon^\star)}{(\gamma+\lambda_1)(\gamma+\lambda_0)}\left(\gamma\lambda_0(1-s+m\phi_1) + \gamma\lambda_1(s-m\phi_1) + \lambda_1\lambda_0\right) < 0. \end{aligned}$$

A.2 Proof of Proposition 2

Existence and uniqueness of a solution to $\Delta(\varepsilon^\star \mid \phi_1) = 0$. It is clear that $\Delta(\varepsilon^\star \mid \phi_1)$ is continuous. Moreover

$$\begin{aligned}
\frac{\partial \Pi_1}{\partial \varepsilon^\star} &= -\left(\frac{\lambda_v}{m}\theta_0 + \frac{\lambda_e N}{m}\phi_0^{N-1}\right)\gamma\left(1 - s + m(1-\phi_0)\right)\frac{1 - F(\varepsilon^\star)}{(r+\gamma+\chi_0)(r+\lambda_0)} \\
\frac{\partial \Pi_0}{\partial \varepsilon^\star} &= \left(\frac{\lambda_v}{m}\theta_1 + \frac{\lambda_e N}{m}\phi_1^{N-1}\right)\gamma\left(s - m\phi_1\right)\frac{F(\varepsilon^\star)}{(r+\gamma+\chi_1)(r+\lambda_1)} \\
\frac{\partial \mathcal{L}}{\partial \varepsilon^\star} &= -1 + \frac{\gamma(1-F(\varepsilon^\star))}{r+\gamma+\chi_0} + \frac{\gamma F(\varepsilon^\star)}{r+\gamma+\chi_1} < -\frac{r}{r+\gamma},
\end{aligned}$$

which implies that $\partial\Delta/\partial\varepsilon^\star < -r/(r+\gamma) < 0$ and bounded away from zero. Hence, $\lim_{\varepsilon^\star\to-\infty}\Delta(\varepsilon^\star \mid \phi_1) = +\infty$ and $\lim_{\varepsilon^\star\to+\infty}\Delta(\varepsilon^\star \mid \phi_1) = -\infty$. It thus follows that the equation $\Delta(\varepsilon^\star \mid \phi_1) = 0$ has a unique solution, which we denote by $\mathcal{E}(\phi_1)$.

$\mathcal{E}(\phi_1)$ is bounded and continuous. Note that $\Delta(\varepsilon^\star \mid \phi_1)$ can be bounded above and below by functions of $\varepsilon^\star$ only. Namely, $\Delta(\varepsilon^\star \mid \phi_1)$ is bounded above by

$$\overline{\pi}\int_{\varepsilon^\star}^{+\infty}\left(\varepsilon - \varepsilon^\star\right)\, dF(\varepsilon) + \left(\frac{\gamma}{r+\gamma} - \underline{\pi}\right)\int_{-\infty}^{\varepsilon^\star}\left(\varepsilon^\star - \varepsilon\right)\, dF(\varepsilon) - \varepsilon^\star + \varepsilon_D \tag{A1}$$

where

$$\overline{\pi} = \frac{\lambda_v\theta_0 + \lambda_e N}{m}\frac{(1-s+m)}{r+\gamma} \text{ and } \underline{\pi} = \frac{\lambda_v\theta_1}{(r+\gamma+\lambda_v(1-\theta_1)+\lambda_e)(\gamma+\lambda_v+\lambda_e)}\frac{\gamma(s-m)}{m}$$

Clearly, the $\varepsilon^\star$ that makes (A1) equal to zero is an upper bound for $\mathcal{E}(\,\cdot\,)$. A lower bound is obtained similarly. The usual sequence-based proof of continuity follows.

The re-statement of dealers' optimality condition. Consider the case $\phi_1 = 0$, which implies that $\Delta(\varepsilon^\star \mid \phi_1) \leq 0$. Since $\Delta(\varepsilon^\star \mid \phi_1)$ is strictly decreasing in $\varepsilon^\star$, this means that $\varepsilon^\star$ has to be greater than $\mathcal{E}(\phi_1)$. A similar reasoning applies to the case $\phi_1 = 1$.

For all m small enough $\mathcal{E}(\phi_1) \in (\underline{\varepsilon}, \overline{\varepsilon})$ for all $\phi_1 \in [0,1]$. We show that, as long as m is small enough, $\Delta(\underline{\varepsilon} \mid \phi_1) > 0 > \Delta(\overline{\varepsilon} \mid \phi_1)$ for all $\phi_1 \in [0,1]$.

To show that $\Delta(\underline{\varepsilon} \mid \phi_1) > 0$ for all m small enough, we first note that $\Pi_0(\underline{\varepsilon} \mid \phi_1) = 0$ so that $\Delta(\underline{\varepsilon} \mid \phi_1) = \Pi_1(\underline{\varepsilon} \mid \phi_1) + \mathcal{L}(\underline{\varepsilon} \mid \phi_1) + \varepsilon_D$. We obtain a lower bound for $\Pi_1(\underline{\varepsilon} \mid \phi_1)$ by evaluating it at $\phi_0 = 0$ everywhere except in $1 - s + m(1-\phi_0)$, which we evaluate at $\phi_0 = 1$:

$$\Pi_1(\underline{\varepsilon} \mid \phi_1) \geq \frac{\lambda_v\theta_0}{m}\gamma(1-s) \times \int_{\underline{\varepsilon}}^{\infty}\frac{\varepsilon - \underline{\varepsilon}}{(r+\gamma+\lambda_v(1-\theta_0)+\lambda_e)(\gamma+\lambda_v+\lambda_e)}\, dF(\varepsilon).$$

Notice that this lower bound goes to infinity as $m \to 0$. We obtain a lower bound for $\mathcal{L}(\underline{\varepsilon} \,|\, \phi_1)$ by setting $\chi_0 = 0$:

$$\mathcal{L}(\underline{\varepsilon} \,|\, \phi_1) \geq -\underline{\varepsilon} - \frac{\gamma}{r+\gamma} \int_{\underline{\varepsilon}}^{\infty} \left(\varepsilon - \underline{\varepsilon}\right) \, dF(\varepsilon) = -\frac{r}{r+\gamma}\underline{\varepsilon} - \frac{\gamma}{r+\gamma}\mathbb{E}\left[\varepsilon\right].$$

Notice that this lower bound does not depend on m. Taken together, we obtain that the sum of these two lower bounds does not depend on ϕ_1 and goes to infinity as m goes to zero. Hence, as long as m is small enough, $\Delta(\underline{\varepsilon} \,|\, \phi_1) > 0$ is strictly positive for all $\phi_1 \in [0,1]$.

Similarly, to show that $\Delta(\overline{\varepsilon} \,|\, \phi_1) < 0$ for all m small enough, we first note that $\Pi_1(\overline{\varepsilon} \,|\, \phi_1) = 0$ so that $\Delta(\overline{\varepsilon} \,|\, \phi_1) = \mathcal{L}(\overline{\varepsilon} \,|\, \phi_1) - \Pi_0(\overline{\varepsilon} \,|\, \phi_1) + \varepsilon_D$. We obtain an upper bound for $-\Pi_0(\overline{\varepsilon} \,|\, \phi_1)$ by evaluating it at $\phi_1 = 0$ everywhere except in $s - m\phi_1$, which is evaluated at $\phi_1 = 1$:

$$-\Pi_0(\overline{\varepsilon} \,|\, \phi_1) \leq -\frac{\lambda_v \theta_1}{m}\gamma(s-m) \times \int_{-\infty}^{\overline{\varepsilon}} \frac{\overline{\varepsilon} - \varepsilon}{(r+\gamma+\lambda_v(1-\theta_1)+\lambda_e)(\gamma+\lambda_v+\lambda_e)} \, dF(\varepsilon).$$

Notice that this upper bound goes to $-\infty$ as $m \to 0$. We can find an upper bound for $\mathcal{L}(\overline{\varepsilon} \,|\, \phi_1)$ by evaluating $\mathcal{L}(\varepsilon^\star \,|\, \phi_1)$ at $\varepsilon^\star = \overline{\varepsilon}$ and $\lambda_v = \lambda_e = 0$:

$$\mathcal{L}(\overline{\varepsilon} \,|\, \phi_1) \leq -\overline{\varepsilon} + \frac{\gamma}{r+\gamma} \int_{-\infty}^{\overline{\varepsilon}} \left(\overline{\varepsilon} - \varepsilon\right) \, dF(\varepsilon) = -\frac{r}{r+\gamma}\overline{\varepsilon} - \frac{\gamma}{r+\gamma}\mathbb{E}\left[\varepsilon\right].$$

Notice that this upper bound does not depend on m. Taken together, we obtain that the sum of these two upper bounds does not depend on ϕ_1 and goes to minus infinity as m goes to zero. Hence, as long as m is small enough, $\Delta(\overline{\varepsilon} \,|\, \phi_1)$ is strictly negative for all $\phi_1 \in [0,1]$.

$\mathcal{E}(\phi_1)$ is strictly decreasing as long as m is small enough. Multiplying both sides of the dealer's optimality condition by m, it is obvious that $\mathcal{E}(\phi_1)$ solves $m\Delta(\varepsilon^\star \,|\, \phi_1) = 0$. Let us calculate the partial derivative of the $m\Delta(\varepsilon^\star \,|\, \phi_1)$ with respect to ϕ_1. Viewing Π_1 as a function of $(\lambda_0, \chi_0, \phi_0^{N-1}, 1-\phi_0)$ and each of these parameters as functions of ϕ_1, we can bound the first term of the partial derivative as follows:

$$\begin{aligned} m\frac{\partial \Pi_1}{\partial \phi_1} &\leq m\frac{\partial \Pi_1}{\partial \lambda_0}\frac{\partial \lambda_0}{\partial \phi_1} + m\frac{\partial \Pi_1}{\partial(1-\phi_0)}\frac{\partial(1-\phi_0)}{\partial \phi_1} \\ &\leq \left(-\frac{\lambda_v^2\theta_0\gamma(1-s)}{(r+\gamma+\lambda_e+\lambda_v)(\gamma+\lambda_e+\lambda_v)^2} + m\frac{\lambda_v\theta_0+\lambda_e N}{r+\gamma}\right)\int_{\varepsilon^\star}^{\infty}(\varepsilon-\varepsilon^\star)\,dF(\varepsilon). \end{aligned}$$

Proceeding similarly we obtain that:

$$-m\frac{\partial \Pi_0}{\partial \phi_1} \leq \left(-\frac{\lambda_v^2\theta_1\gamma(s-m)}{(r+\gamma+\lambda_e+\lambda_v)(\gamma+\lambda_e+\lambda_v)^2} + m\frac{\lambda_v\theta_1+\lambda_e N}{r+\gamma}\right)\int_{-\infty}^{\varepsilon^\star}(\varepsilon^\star-\varepsilon)\,dF(\varepsilon).$$

Finally

$$m\frac{\partial \mathcal{L}}{\partial \phi_1} \leq \frac{\gamma m}{(r+\gamma+\chi_1)^2}\left(-\frac{\partial \chi_1}{\partial \phi_1}\right)\int_{-\infty}^{\varepsilon^\star}\left(\varepsilon^\star-\varepsilon\right)\,dF(\varepsilon)+\frac{\gamma m}{(r+\gamma+\chi_0)^2}\left(\frac{\partial \chi_0}{\partial \phi_1}\right)\int_{\varepsilon^\star}^{\infty}\left(\varepsilon-\varepsilon^\star\right)\,dF(\varepsilon).$$

Taking the derivative with respect to ϕ_1, we find that both $-\partial\chi_1/\partial\phi_1$ and $\partial\chi_0/\partial\phi_1$ are less than $\lambda_v+\lambda_e N(N-1)$. Hence, we have that:

$$m\frac{\partial \mathcal{L}}{\partial \phi_1} \leq \frac{\gamma m(\lambda_v+\lambda_e N(N-1))}{(r+\gamma)^2}\left(\int_{-\infty}^{\varepsilon^\star}\left(\varepsilon^\star-\varepsilon\right)\,dF(\varepsilon)+\int_{\varepsilon^\star}^{\infty}\left(\varepsilon-\varepsilon^\star\right)\,dF(\varepsilon)\right).$$

Taken together, we obtain that $m\partial\Delta/\partial\phi_1 \leq H(\varepsilon^\star, m)$, for some continuous function of $(\varepsilon^\star, m)$ such that $H(\varepsilon^\star, 0) < 0$ for all $\varepsilon^\star \in [\underline{\varepsilon}, \overline{\varepsilon}]$. It thus follows by continuity that, as long as m is small enough, $H(\varepsilon^\star, m)$ is strictly negative for all $\varepsilon^\star \in [\underline{\varepsilon}, \overline{\varepsilon}]$.

A.3 Proof of Proposition 3

We first establish that the equilibrium set is comprised of all $(\phi_1, \varepsilon^\star)$ such that $\phi_1 = \Phi_1 \circ \mathcal{E}(\phi_1)$ and:

$$\varepsilon^\star \begin{cases} \in [\mathcal{E}(\phi_1), \underline{\varepsilon}] & \text{if } \phi_1 = 0 \\ = \mathcal{E}(\phi_1) & \text{if } \phi_1 \in (0,1) \\ \in [\overline{\varepsilon}, \mathcal{E}(\phi_1)] & \text{if } \phi_1 = 1 \end{cases}$$

Let us show that any equilibrium, $(\phi_1, \varepsilon^\star)$, must be included in the set defined above. *If* $\phi_1 \in (0,1)$, the optimality condition of dealers implies that $\varepsilon^\star = \mathcal{E}(\phi_1)$. Market-clearing implies that $\phi_1 = \Phi_1(\varepsilon^\star)$, and we are done. *If* $\phi_1 = 0$, then Proposition 1 implies that $\varepsilon^\star \leq \underline{\varepsilon}$, and Proposition 2 implies that $\varepsilon^\star \geq \mathcal{E}(0)$. Taken together, this implies that $\varepsilon^\star \in [\mathcal{E}(0), \underline{\varepsilon}]$. Since $\mathcal{E}(0) \leq \underline{\varepsilon}$, it follows from Proposition 1 that $\Phi_1 \circ \mathcal{E}(0) = 0$, and we are done. The $\phi_1 = 1$ case is analogous.

Next, we show the reverse inclusion: any $(\phi_1, \varepsilon^\star)$ belonging to the set defined above is an equilibrium. *If* $\phi_1 \in (0,1)$, then it is clear that both the market-clearing and the dealer's optimality condition hold. *If* $\phi_1 = 0$, then the market-clearing condition holds because $\varepsilon^\star \leq \underline{\varepsilon}$. Likewise, the optimality condition of dealers holds because $\varepsilon^\star \geq \mathcal{E}(\phi_1)$. The $\phi_1 = 1$ case is analogous.

Then, we show that an equilibrium exists. First, given that we have established in Proposition 1 and Proposition 2 that both $\Phi_1(\cdot)$ and $\mathcal{E}(\cdot)$ are continuous, the Intermediate Value Theorem implies that the equation $\phi_1 = \Phi_1 \circ \mathcal{E}(\phi_1)$ has at least one solution. Clearly, the solution of this fixed point problem, together with $\varepsilon^\star = \mathcal{E}(\phi_1)$, belongs to the set defined in the proposition.

Proposition 1 shows that $\Phi_1(\cdot)$ is increasing in $\varepsilon^\star \in (\underline{\varepsilon}, \overline{\varepsilon})$, while Proposition 2 shows that, as long as m is small enough, $\mathcal{E}(\phi_1) \in (\underline{\varepsilon}, \overline{\varepsilon})$ and is strictly decreasing. This implies that the function $\phi_1 \mapsto \Phi_1 \circ \mathcal{E}(\phi_1) - \phi_1$ is strictly decreasing. It follows that the equilibrium is unique. Moreover, since $\mathcal{E}(\phi_1) \in (\underline{\varepsilon}, \overline{\varepsilon})$, it follows from Proposition 1 that $\phi_1 = \Phi_1 \circ \mathcal{E}(\phi_1) \in (0,1)$.