Frictional Intermediation, Inventory Hedging, and the Rise of Portfolio Trading in the Corporate Bond Market

Jessica Shi Li[†]

University of Chicago Booth School of Business

This Version: March 23, 2022

Abstract

The rapid rise of corporate bond portfolio trading since the end of 2017 has attracted attention from practitioners and regulators alike. I show that inventory hedging explains the recent meteoric rise of corporate bond portfolio trading, likely aided by the recent proliferation of credit index derivatives. To formally study inventory hedging and its implications for dealer liquidity provision and price formation, I build a simple dynamic search model with frictional intermediation and inventory hedging. The model predicts that inventory hedging improves liquidity and transaction costs, and the liquidity benefit is more pronounced when dealers face higher inventory frictions or constraints. Given that portfolio trading is driven by inventory hedging, I show that consistent with model predictions, portfolio trading is associated with lower transaction costs, and is more beneficial for bonds with higher ex-ante inventory frictions and when the dealer sector is constrained. I also document that portfolio trading reduces transaction costs of similar voice trades, and attribute the cross venue effect to search and information spillover between trading protocols.

Keywords: OTC, Liquidity, Dealer, Search, Inventory, Hedging, Corporate Bonds, Portfolio Trading

JEL Codes: G12, G23, G24, G28

[†]University of Chicago Booth School of Business. Email: jessicali@chicagobooth.edu. This research was funded in part by the John and Serena Liew Fellowship Fund at the Fama-Miller Center for Research in Finance, University of Chicago Booth School of Business. I thank Doug Diamond, Zhiguo He, Anil Kashyap, Ralph Koijen, Stefan Nagel, Wenxin Du, Yueran Ma, and participants at the Chicago Booth Asset Pricing Working Group for helpful discussions. I thank Yao Zeng for helpful insights on corporate bond ETFs.

1 Introduction

A plethora of asset classes, including fixed income securities, commodities, and various derivatives, are traded in over-the-counter (OTC) markets intermediated by dealers. One such example is the corporate bond market, which is of particular interest given its large size and importance in firm financing. The bilateral and sequential nature of corporate bond trading leads to low liquidity and high transaction costs. The disruption in March 2020 during the COVID-19 crisis further highlights liquidity-induced systemic risk, making bond market liquidity front and center in policy debates.

Compared to the equity market, the corporate bond market seems slow to innovate. Although corporate bond electronic trading has developed in the past decade, trading on electronic platforms remains concentrated in high-quality bonds and small trade sizes (e.g., Hendershott and Madhavan (2015), Kozora et al. (2020), O'Hara and Zhou (2021b)). Liquidity has shown little improvement in the institutional segment of the market as well as for high-yield (HY) bonds. However, a recent innovation is making significant inroads into the corporate bond market, challenging to transform trading and liquidity in this market. This innovation is corporate bond portfolio trading, or trading of a portfolio of bonds at the same time pursuant to a single order. Figure 1 shows the remarkable increase in corporate bond portfolio trading, which was nearly non-existent prior to the end of 2017, but has experienced fast growth since then, reaching above 5% of total investor-dealer trading volume in the U.S. corporate bond market as of June 2021.

Figure 1: The Rise of Portfolio Trading in the U.S. Corporate Bond Market

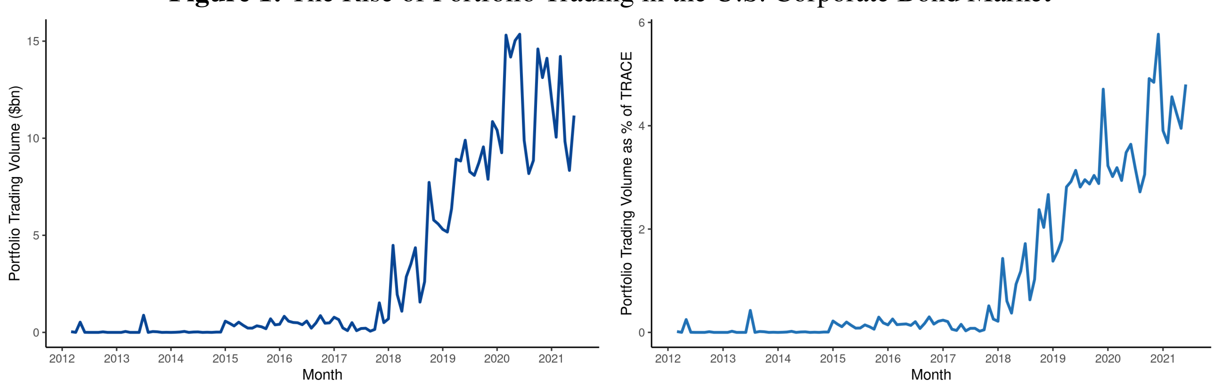

Note: The left panel shows the monthly portfolio trading volume (in $bn) between March 2012 and June 2021. The right panel shows the monthly portfolio trading volume as a percentage of total investor-dealer trading volume according to TRACE during the same period. I define a portfolio trade as a basket of at least 30 unique individual bonds traded at the exact same time between a dealer and a client. Identification of portfolio trades in the TRACE dataset is explained in Section 2.2. Data are compiled from Enhanced TRACE.

An important question is what caused such rapid rise of corporate bond portfolio trading. The initial rise of portfolio trading in late-2017 to early-2018 did not coincide with any related regulatory change. So, if not regulations, what led to the emergence of portfolio trading?

A demand-side explanation is that the growth in passive index funds has increased the demand for trading multiple bonds all at once. However, if the demand-side is the main driving force, there should be little observed differences between portfolio trading and voice trading in terms of dealers' market-making. On the contrary, I find that while the transaction costs of voice trades are sensitive to inventory shocks, the transaction costs of portfolio trades do not respond to such shocks. Moreover, I find that the average turnover in portfolio-trade bonds is about 50% slower than that in comparable voice-trade bonds. In other words, dealers' inventory behavior and management appear to significantly differ from voice trading to portfolio trading. These empirical observations suggest that the demand-side story alone is inadequate. Thus, it is important to study the supply side, i.e., intermediation by dealers.

Could these empirical observations be explained by dealers' reliance on the ETF creation and redemption mechanism, given the growth of corporate bond ETFs over the last decade? A dealer who has purchased a portfolio of bonds through portfolio trading may deliver those bonds to the ETF sponsor, in exchange for an ETF share that is highly liquid. Similarly, a dealer who has sold a portfolio of bonds can source those bonds from the ETF sponsor by redeeming an ETF share. However, if the ETF mechanism plays a significant role, we would expect that bonds traded through portfolio trading are largely ETF basket bonds. I follow Shim and Todorov (2021) and Reilly (2021) to infer realized ETF baskets across corporate bond ETFs with in-kind creation and redemption mechanism, and find that the overlap between portfolio-trade bonds and ETF-basket bonds is only 20-30%. Moreover, portfolio trading does not seem correlated with ETF creation and redemption activity, further suggesting that the ETF mechanism is insignificant.

Instead, I argue that the rise of corporate bond portfolio trading is driven by dealer inventory hedging. The empirical observations above are consistent with this hypothesis. Testing the hypothesis directly is challenging, as dealers are not required to disclose information on their hedging activities, and their hedging behavior is unobserved. However, since inventory hedging leads to differences in dealers' intermediating behavior, by observing such differences in the data, I can reasonably conclude that portfolio trading is related to inventory hedging. First, if a trade is unhedged or results in increased inventory exposure, its transaction cost should be increasing in contemporary inventory shocks. This is because positive inventory shocks increase dealers' marginal cost of inventory, resulting in lower propensity to provide liquidity and higher transaction costs. The finding that portfolio-trade bonds lack sensitivity to inventory shocks is thus consistent with the inventory hedging hypothesis. Second, if portfolio trades are associated with better ability to hedge inventories, dealers are less incentivized to search for counterparties and enter into offsetting trades, resulting in slower inventory turnover for portfolio-trade bonds. The finding that portfolio-trade bonds stay in inventories longer supports this, and rules out the possibility that portfolio trades are pre-arranged and hence do not respond to inventory shocks. Furthermore, the slower turnover of portfolio-trade bonds is inconsistent with the information explanation. If portfolio-trade bonds' insensitivity to inventory shocks is a result of lower information sensitivity at the portfolio-level and reduced adverse selection, these bonds should turn over faster due to improved liquidity. Moreover, using realized equity volatilities to proxy private information, I find that transaction cost differentials between portfolio trading and voice trading do not respond to the amount of private information.

In fact, the rise of portfolio trading coincides with the boom in credit index derivatives, which make good candidates for hedging portfolios of bonds. The post-GFC reform requires that all standardized OTC derivatives must be cleared with central counterparties by the end of 2012. As a result, a highly liquid market for standardized and centrally cleared credit index derivatives developed overtime. For example, Markit iBoxx Standardized Total Return Swaps (TRS) started trading in 2015 and have experienced rapid growth in trading volumes ever since (Figure A.1). Similarly, not only have corporate bond ETFs grown significantly in assets, but a highly active market has formed for exchange-traded corporate bond ETF options (Figure A.2). Thus, in contrast to single-name CDS contracts which are typically traded bilaterally and often illiquid, credit index derivatives provide a liquid and more effective solution to hedging bond inventories.[1]

To provide intuition and illustrate how inventory hedging works and affects liquidity provision and transaction costs, I then build a stylized model featuring dynamic search and frictional intermediation, built upon Duffie et al. (2005). I extend their model in two dimensions. First, I assume that the inter-dealer market trading is bilateral and frictional, with similar search and bargaining as those in investor-dealer trading, consistent with recent empirical evidence (e.g., Di Maggio et al. (2017), Hollifield et al. (2017)). Thus, dealers often encounter costly inventory imbalances, given these frictions. Second, I introduce a hedging technology that reduces the inventory holding cost of hedged positions. However, dealers are subject to i.i.d. shocks to their access to this hedging technology. The accessibility of the hedging technology affects dealers' intermediating decisions.

[1]A key difference between portfolio trading and traditional voice trading is that in portfolio trading, a portfolio of bonds are transacted at the same time. Thus, dealers can arrange hedging for the resulting inventories at the time of portfolio trades. In contrast, dealers trade individual bonds throughout the day through voice trading, and thus at the time of each transaction, dealers do not yet know what their end-of-day positions are and how well they can be hedged. Furthermore, due to specialization, dealer inventory management is decentralized across traders within the dealer firm (e.g., Naik and Yadav (2003a), Nagel (2012), Cespa and Foucault (2014)), further complicating hedging of voice trades.

The economic meaning of inventory hedging is worth discussing. In the model, the inventory technology is exogenous and reduces dealers' inventory holding cost. Inventory holding cost arises when a dealer's inventory positions deviate from its desired holdings, resulting in inventory imbalance (e.g., Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981)). In practice, capital is allocated to specialized traders within a dealer firm, and thus market-making tends to be decentralized (e.g., Naik and Yadav (2003a), Nagel (2012), Cespa and Foucault (2014)). Traders are typically subject to risk budgets, determined by internal risk management parameters of the dealer. Thus, inventory imbalances are often costly due to the Value-at-Risk (VaR) constraint, as well as increased career risk of traders and heightened compliance and monitoring. Inventory hedging lowers inventory costs, by relaxing these constraints through risk reduction.

The model shows that when the inventory hedging technology is accessible, dealers are more willing to intermediate, resulting in lower bid-ask spreads and higher liquidity. Furthermore, hedging accessibility interacts with dealer inventory frictions. The accessibility of the hedging technology is more beneficial when inventory frictions are high. Given that corporate bond portfolio trading is driven by dealer inventory hedging, the model thus yields two key testable predictions in the context of portfolio trading. First, portfolio trades are associated with lower transaction costs, compared to similar voice trades. Second, the transaction cost advantage of portfolio trading is more pronounced for bonds with higher ex-ante inventory frictions (e.g., bonds with lower credit quality, longer maturities and smaller issue sizes), and when the dealer sector is more constrained so that inventory frictions are elevated.

I then empirically test these predictions. An ideal experiment would be to randomly assign dealers into two groups, so that one (treatment group) trades corporate bonds through portfolio trading, while the other (control group) trades the same bonds via traditional voice trading. Such experiment is not feasible in reality. However, since portfolio trading and voice trading co-exist, and a given bond may be transacted through both portfolio trading and voice trading by a given dealer or on a given day, portfolio-trade bonds can be compared to voice-trade bonds, while controlling for many confounding factors such as differences across bond issues, dealers, trade sizes and time. I show that, consistent with model predictions, portfolio trading is associated with lower transaction costs relative to voice trading, at both the portfolio level and individual bond level. Moreover, portfolio trading tends to provide more benefit to bonds with lower credit ratings, longer remaining maturities and smaller issue sizes. Portfolio trading also yields higher liquidity benefit when the dealer sector is more constrained, as proxied by high TED rate. I further document an interesting cross venue effect, that is, portfolio trading also reduces the transaction costs of similar voice trades by the dealer. I attribute this effect to search and information spillover from portfolio trading to voice trading.

1.1 Related Literature

To the best of my knowledge, this paper is the first to study portfolio trading in the corporate bond market. There exist a few studies on equity portfolio trading (e.g., Stoll (1988), Hill and Jones (1988)). The corporate bond market is significantly different from the equity market due to its decentralized nature. The finding that corporate bond portfolio trading is driven by dealer inventory hedging distinguishes this paper from the literature on equity portfolio trading. This paper is related to the literature on dealers' market-making and inventory management. Studies by Friewald and Nagler (2019), Goldstein and Hotchkiss (2020), and He et al. (2021) find that dealer inventory concerns feature prominently in OTC markets. While the existing literature has focused on dealers setting prices (e.g., Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981)) and varying search intensities (e.g., Goldstein and Hotchkiss (2020)) to manage their inventories, I show that inventory hedging plays a significant role in dealer inventory management. Although a few empirical studies have chronicled inventory hedging in several OTC markets (e.g., Naik and Yadav (2003b), Fleming and Rosenberg (2008), Siriwardane (2019), Chen et al. (2021)), a key contribution of this paper is to formally study dealer inventory hedging.

This paper also contributes to the theoretical OTC search literature.[2] In contrast to the seminal search-theoretic models of OTC markets (e.g., Duffie et al. (2005), Lagos and Rocheteau (2009)), this paper models a frictional inter-dealer market subject to similar OTC search frictions found in investor-dealer trading, and thus dealers experience inventory shocks from time to time. Recent studies show that inter-dealer markets are often bilateral and frictional (e.g., Green et al. (2007), Li and Schürhoff (2019), Hollifield et al. (2017), Bessembinder et al. (2020)). Instead of modeling a fully decentralized trading market (e.g., Afonso and Lagos (2011), Üslü (2019), Farboodi et al. (2019, 2021), Hugonnier et al. (2021)), I take the two-tier market structure as given, that is, investors can only trade with dealers while dealers can also trade with each other in an inter-dealer market. A related paper is Hugonnier et al. (2020) who also model a frictional inter-dealer market featuring search and bargaining friction and dealers with heterogeneous inventory costs. However, while this paper focuses on inventory hedging and its role in dealer inventory management, Hugonnier et al. (2020) do not consider hedging and instead emphasize intermediation chains as a result of frictional inter-dealer trading. Their paper models

[2] See Weill (2021) for a survey of the literature.

time-invariant heterogeneity across dealers in terms of holding cost. I assume that dealers are ex-ante symmetric when meeting investors but their access to the hedging technology is stochastic. This makes the problem easy to solve and eliminates the issue that dealers with different inventory cost may have varying bargaining power against investors. Yang and Zeng (2021) also model a two-tier, but fully decentralized market. They study dealer liquidity provision, assuming dealers can hold more than one unit of asset, and uncover multiple equilibria. In this paper, I instead focus on the impact of intermediary constraints and search on liquidity, and derive a unique steady state equilibrium.

This paper is also related to the literature that studies the impact of intermediary constraints on liquidity (e.g., Gromb and Vayanos (2002), Brunnermeier and Pedersen (2008), Nagel (2012)), as well as the literature on intermediary asset pricing (e.g., He and Krishnamurthy (2012, 2013), He et al. (2017)), which studies the impact of intermediary frictions on asset prices and risk premia. However, a key insight of this paper is the interdependence of OTC search frictions and dealer inventory frictions. That is, the search frictions inherent in the OTC markets amplify the effect of dealer inventory frictions, and vice versa. Furthermore, this paper relates to the literature on financial regulations and their impact. For example, Bao et al. (2018), Bessembinder et al. (2018), Dick-Nielsen and Rossi (2018) examine the impact of post-GFC regulatory restrictions (such as the Volcker Rule) on dealer capital commitment and liquidity provision. Du et al. (2017), Fleckenstein and Longstaff (2020), Du and Schreger (2021) study bank balance sheet space constraints. This paper also adds to the literature that studies financial innovations in the market structure of OTC markets (e.g., Hendershott and Madhavan (2015), Riggs et al. (2020), O'Hara and Zhou (2021b)).

The remainder of the paper is organized as follows. Section 2 documents some new stylized facts on corporate bond portfolio trading and provides an overview of the institutional background as well as description of the data. Section 3 tests that portfolio trading is associated with dealer inventory hedging. Section 4 lays out a conceptual framework of dealer inventory hedging and its impact, by developing a simple model featuring dynamic search and frictional intermediation. Section 5 tests model predictions and studies the effect of portfolio trading on transaction costs. Section 6 concludes.

2 Institutional Background and New Facts

In this section, I provide relevant institutional background and present some new stylized facts on corporate bond portfolio trading. I also describe data used for the empirical analyses in this paper.

2.1 Corporate Bond Transactions

Corporate bonds in the U.S. have traditionally traded OTC with dealers serving as intermediaries. In investor-dealer “voice" trades, dealers disseminate (e.g., through emails or Bloomberg messages) bid and ask quotes for bonds and may also indicate whether they are a buyer or seller of a bond and quantity bid or offered. Given these quotes, investors strategically contact dealers based on their networks of dealers and estimated execution quality. The standard voice trading protocol requires the investor to contact dealers sequentially. Historically, negotiations occurred over the phone, and hence the name voice trading. With advances in technology, voice trades are now more likely to occur via instant messengers such as Bloomberg instant chats. These bilateral voice trades feature a two-tier market structure. That is, investors must trade with dealers and cannot transact with other investors directly, while dealers can also trade with each other in the inter-dealer market, facilitated by specialized voice brokers. In a bilateral voice trade, both investors and dealers experience significant search and bargaining frictions, leading to low execution certainty and unfavorable prices (Duffie et al. (2005, 2007)).

I obtain U.S. corporate bond transaction data between March 1, 2012 and December 31, 2018 from the academic version of TRACE, provided by the Financial Industry Regulatory Authority (FINRA). There are around 121 million transactions involving more than 85,000 unique CUSIPs. The Academic TRACE contains trade-level information on U.S. corporate bond transactions, including CUSIPs, trade prices, uncapped transaction sizes, trade time-stamps, dealer buy/sell indicators, dealer identifiers, counterparty information, trade status and trading market indicators. A key advantage of the Academic TRACE is that the dataset contains anonymized dealer codes identifying the transacting dealers on each trade. This allows for exploitation of dealer-level differences. My main estimation sample covers the period between January 1, 2018 and December 31, 2018.[3]

The estimation sample covers 22.7 million transactions involving 41,308 unique CUSIPs. I follow the literature and apply a standard filtering procedure (e.g., Dick-Nielsen and Poulsen (2019)) to clean the Academic TRACE dataset. I then merge the cleaned TRACE dataset with bond characteristics and credit ratings data obtained from the Mergent Fixed Income Securities Database (FISD). I drop bonds whose CUSIPs have no match in the FISD dataset. Following the literature (e.g., Friewald and Nagler (2019)), I also restrict my sample to U.S. corporate debentures that have fixed coupons and are not convertible, asset-backed, exchangeable, privately placed, perpetual, or preferred securities. I further exclude any secured lease obligations and bonds quoted in foreign currencies. I then delete observations whose settlement dates occur on or after the bond

[3]The January 1, 2018 start date roughly corresponds to the initial rise of corporate bond portfolio trading. The December 31, 2018 end date is due to a 36-month lag of academic TRACE.

maturity dates and observations with non-positive reported prices. Since the focus of this paper is on the effect of inventory hedging on investor-dealer trading liquidity, and because portfolio trading caters to buy-side investors, I exclude inter-dealer trades and primary market transactions from the estimation sample. To facilitate the comparison of portfolio trades to traditional voice trades, I also remove electronic trades from my sample by deleting any trade with an ATS flag. This yields a sample containing 6,369,792 transactions and 11,398 unique issue CUSIPs.

2.1.1 Transaction Cost

Following Hendershott and Madhavan (2015) and O'Hara and Zhou (2021a,b), I use the most recent inter-dealer trade in a bond as the benchmark price and define transaction cost as:

$$Cost_k = \ln\left(\frac{Trade\ Price_k}{Benchmark\ Price_k}\right) \times Trade\ Sign_k \tag{1}$$

where $Trade\ Price_k$ is the transaction price for trade k, $Benchmark\ Price_k$ is the transaction price of the most recent inter-dealer trade in the same bond, and $Trade\ Sign_k$ indicates the direction of the trade and takes on the value of $+1$ for dealer sell (investor buy) and -1 for dealer buy (investor sell).

Hendershott and Madhavan (2015) compare this measure of corporate bond transaction cost to several other measures used in the literature. O'Hara and Zhou (2021a) also address the advantages of using this measure. A commonly used illiquidity measure is the Amihud (2002) measure, defined as the average ratio of the daily absolute return to the trading volume on the same day. Thus, calculating the Amihud measure on a daily frequency requires at least two intra-day observations (transactions) of a security. Since many corporate bonds trade infrequently, applying the Amihud measure to the corporate bond market has serious drawbacks. To study the corporate bond market, researchers often use the dealers' round-trip spread (or effective bid-ask spread), estimated by subtracting average dealer-buy prices from average dealer-sell prices on a given day (e.g., Hong and Warga (2000)). The round-trip spread measure has the same drawback as the Amihud measure, as it requires both buy and sell trades of a bond on a single day. In the context of portfolio trading, since portfolio trades are often one-side - that is, a bond traded as part of a portfolio trade usually does not have an offsetting trade that is also in a portfolio transaction - calculating the round-trip spreads for portfolio trade bonds and non-portfolio trade bonds separately can be challenging. The transaction cost measure (1) shares similarity with a regression approach to estimate transaction cost (e.g., Schultz (2002) and Bessembinder et al. (2006)), and can be easily computed.

To illustrate the levels and time trends of the overall transaction costs for different trade size buckets and credit ratings, I plot the yearly average transaction costs by trade size for investment-grade (IG) and high-yield (HY) bonds separately, in Figure 2. I follow the literature convention as well as industry practice and divide trade sizes into four categories, including micro (< \$100,000), odd lot (\$100,000 - \$1 million), round lot (\$1 - 5 million) and block ($\geq$ \$5 million). Consistent with Hendershott and Madhavan (2015) and O'Hara and Zhou (2021b), transaction costs of micro trades dropped precipitously in the past decade, benefiting from an increase in the share of electronic trading in retail-sized trades. Transaction costs of odd lots also decreased materially, for both IG and HY bonds. Meanwhile, transaction costs of round lot and block trades remained relatively stable, at around 10 bps for IG bonds and 20 bps for HY bonds.

Figure 2: Transaction Costs by Trade Size

(a) Investment Grade (IG) **(b)** High Yield (HY)

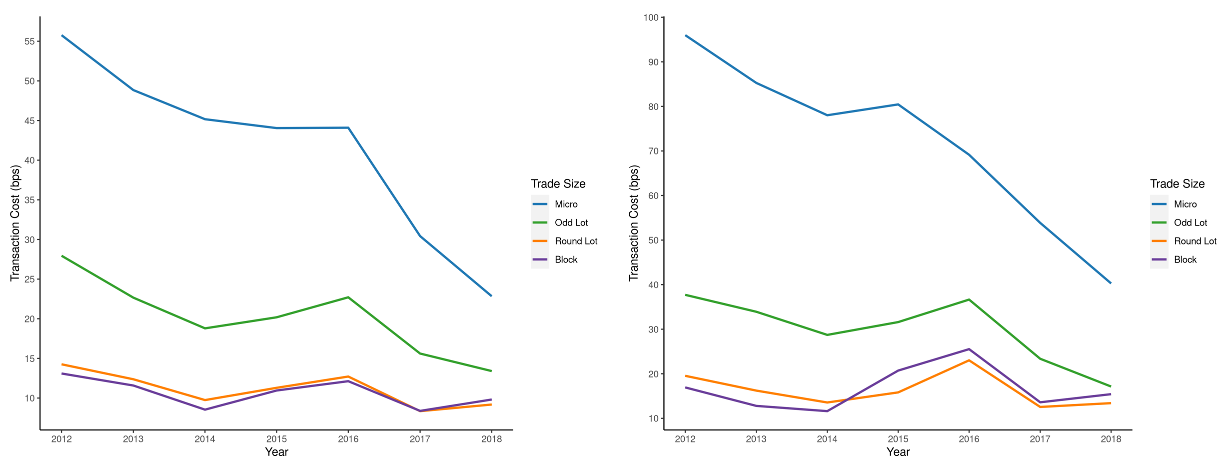

Note: The figure plots the yearly average transaction costs by trade size buckets from 2012 to 2018, for both investment grade (IG) and high yield (HY) bonds.Trade size buckets include micro (< \$100,000), odd lot (\$100,000 - \$1 million), round lot (\$1-\$5 million), and block ($\geq$ \$5 million). Transaction costs are reported in basis points.

2.1.2 Bond Turnover

To compare dealer holding periods across bond transactions, I construct a bond turnover measure defined as the time period between the original transaction and subsequent transactions by the same dealer in the same bond and of the opposite direction, for which the cumulative volume fully offsets the trade volume of the original transaction. This method takes into account the possibility that offsetting trades may be broken up into smaller pieces. For example, a \$5 million dealer-buy trade may be matched to two subsequent dealer-sell trades of \$2 million and \$3 million, respectively. Importantly, since the focus is on bond turnover at individual dealer level, I include all inter-dealer trades by the same dealer that are of the opposite direction to the original transaction. The bond turnover measure proxies, for each trade, how fast a dealer enters into offsetting transactions. The mean turnover in the estimation sample is about 11 days.[4] Turnover in bonds purchased by dealers may be of particular interest, given that dealers typically do not carry significant short positions (e.g., Yang and Zeng (2021)). In this case, the turnover measure proxies how long these bonds stay on dealers' balance sheets.

2.2 Portfolio Trades

A portfolio trade involves a single transaction between two parties (a dealer and an investor) for a basket of corporate bonds, executed on an all-or-none basis pursuant to a single order. Instead of contacting dealers sequentially, an investor sends a list of bonds to a few dealers "in competition", similar to the RFQ protocol in electronic trading. These dealers then respond within a predetermined time frame (e.g., 30 minutes) with prices for the entire portfolio. The investor either accepts the best quote from the dealers or rejects all quotes. If the investor accepts the best quote, the portfolio trade is then executed with the dealer who has provided the best quote. If the investor rejects all quotes, no trade takes place. While a portfolio trade is executed on an all-or-none basis with the basket of securities transacting at the same time, the dealer is required to report the transaction information for each line item in the portfolio trade to TRACE.

Portfolio trades are not flagged in the TRACE dataset. However, since bonds traded as part of a portfolio trade are executed at the same time, I approximately identify portfolio trades as baskets of unique CUSIPs traded at the exact same time by the same dealer with an investor. I also exclude any time-stamp at which a single CUSIP was transacted multiple times. This removes spot trades and other session-based trades that also tend to occur at approximately the same time.[5]

2.2.1 Rise of Portfolio Trading

Figure 3 plots the monthly "portfolio" trading volume from March 2012 to December 2018, where portfolio trades are identified from the data using different number-of-CUSIPs thresholds. Table 1 reports the monthly average number of such portfolio transactions and corresponding notional size, for 2018 (the estimation sample

[4]The mean turnover is consistent with the estimate by Goldstein and Hotchkiss (2020). If excluding trades with an exactly offsetting trade within 15 minutes, the mean bond turnover increases to 14 days. Roundtrip trades within 15 minutes likely proxy agency trades. However, since agency trades cannot be accurately identified within the TRACE dataset and the inclusion of agency trades yield more conservative estimates in my empirical analysis, I keep these roundtrip trades in the estimation sample.

[5]Investment grade (IG) bonds trade at spreads to benchmark Treasury yields, and then converted to dollar prices via "spotting" the corresponding Treasury security. The spotting process is sometimes delayed to a set time of the day (e.g., 3pm in the afternoon).

period), the two years immediately preceding 2018, as well as the pre-2018 period from March 2012 to December 2017. Regardless of the choice of cutoffs, both the number and notional size of identified portfolio transactions each month remained low in magnitude and stable prior to 2018, and increased sharply in 2018. The increase in portfolio transactions with at least 30 line items was particularly stark. Portfolio trades containing at least 30 line items occurred merely 11 times ($284 million notional) a month from March 2012 to December 2017, increasing to 68 times ($3,954 million notional) a month on average in 2018. Furthermore, even for identified portfolio trades with 10 to 30 line items, the growth rate in notional size per month was more than double that of the number of trades. This suggests that the size of these portfolio transactions expanded significantly in 2018.

Figure 3: Monthly Volume of Portfolio Trades by Number of Line Items

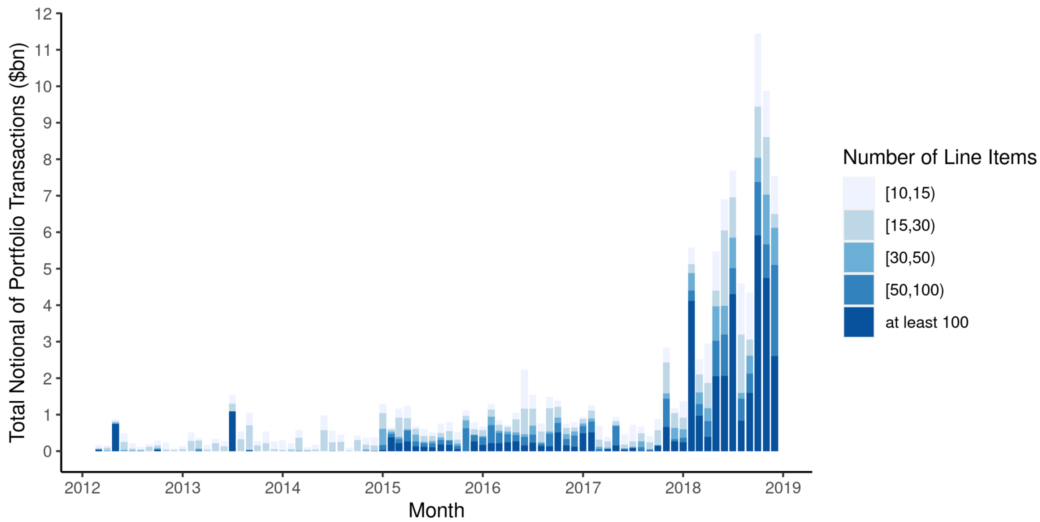

Note: This figure plots monthly corporate bond portfolio trading volume (in $bn) from March 2012 to December 2018, where portfolio trades are broken down to five categories based on the number of line items, including portfolio trades with 10-15 line items, 15-30 line items, 30-50 line items, 50-100 line items and at least 100 line items. Data are compiled using corporate bond transaction data from the Academic TRACE dataset.

Based on the evidence in Table 1, and knowing that the identification of portfolio trades using the above method is not exact, a number-of-CUSIPs cutoff of 30 seems more reasonable. A higher number-of-CUSIPs threshold and the availability of dealer codes in the Academic TRACE mitigate the risk of mis-classification. First, since the execution time-stamps in the TRACE dataset are reported to the second, it would be unlikely for many bonds to be executed at the same second by a single dealer by happenstance, after removing spot trading and other session-based trading from calculation. Second, a cutoff of 30 line items addresses the potential concern that some bonds traded as part of a list trade (such as BWIC/OWIC) may be mis-classified as part

Table 1: Portfolio Trades Summary Statistics

	Pre-2018	2016	2017	2018
Portfolio Transactions per Month				
Obs - No. Line Items [10,15)	125	154	143	308
Obs - No. Line Items [15,30)	51	93	70	143
Obs - No. Line Items [30,50)	5	9	7	24
Obs - No. Line Items [50,100)	4	10	5	22
Obs - No. Line Items $\geq$ 100	2	5	5	22
Notional Size per Month ($mn)				
Notional - No. Line Items [10,15)	197	253	234	1,006
Notional - No. Line Items [15,30)	212	307	262	899
Notional - No. Line Items [30,50)	54	107	80	637
Notional - No. Line Items [50,100)	103	247	173	832
Notional - No. Line Items $\geq$ 100	127	214	203	2,485

Note: This table reports the monthly average number of portfolio trades and corresponding notional size for the pre-2018 period (March 2012 - December 2017), 2016, 2017 and 2018 respectively. Portfolio trades are defined as baskets of unique CUSIPs traded at the exact same time by the same dealer with a customer. Any time-stamp at which a single CUSIP was transacted multiple times is excluded. Various cutoffs for the number of unique CUSIPs are used to identify portfolio transactions; in particular, portfolio trades with 10-15 line items, 15-30 line items, 30-50 line items, 50-100 line items, and at least 100 line items.

of a portfolio trade. List trading has long been in use in the fixed income market. Like portfolio trading, list trading also involves pricing and trading of lists of securities instead of individual items one by one. However, while portfolio trading is executed with one dealer all at once, list trading typically involves partial execution with many different dealers. In list trades, individual line items are broadcast widely to many dealers who then cherry-pick the bonds that they are willing to trade, and thus many bonds in a list do not transact at all. As a result, the number of CUSIPs transacted by the same dealer at the same time pursuant to a list trade is generally small. The existence of portfolio trades identified in the pre-2018 sample and the significant increase in trade sizes observed in 2018 also suggest that my method picks up some list trades in the pre-2018 period, but misclassification is significantly reduced by using a higher number-of-CUSIPs threshold.[6]

Therefore, for the rest of the paper, I define a portfolio trade as a basket of at least 30 unique CUSIPs transacted at the same time between a dealer and a customer. This threshold is consistent with some recent

[6] Given that TRACE volume stayed relatively stable from March 2012 to December 2018 (see Figure A.4) and that list trading was already a mature trading protocol, potential mis-classification is unlikely to explain the rise of portfolio trading documented in this section. A potential threat is that list trading may have experienced a significant change around the end of 2017. For example, if list trades started to include significantly more line items or if execution of list trades became materially more concentrated among a few dealers, then mis-classification of list trading as portfolio trading would be more likely. However, such change would make list trading substantially similar to portfolio trading, and for the purpose of this paper, I do not differentiate between portfolio trades and “portfolio-like" list trades.

regulatory recommendations,[7] as well as industry conventions. It is worth noting that the estimated portfolio trading volume was not exactly zero before 2018, even with a cutoff of 30 unique CUSIPs, likely because my method picks up some list trades involving large number of bonds and some (although few) bespoke portfolio transactions did take place prior to the end of 2017.

Figure 4: Corporate Bond Portfolio Trading Monthly Volume

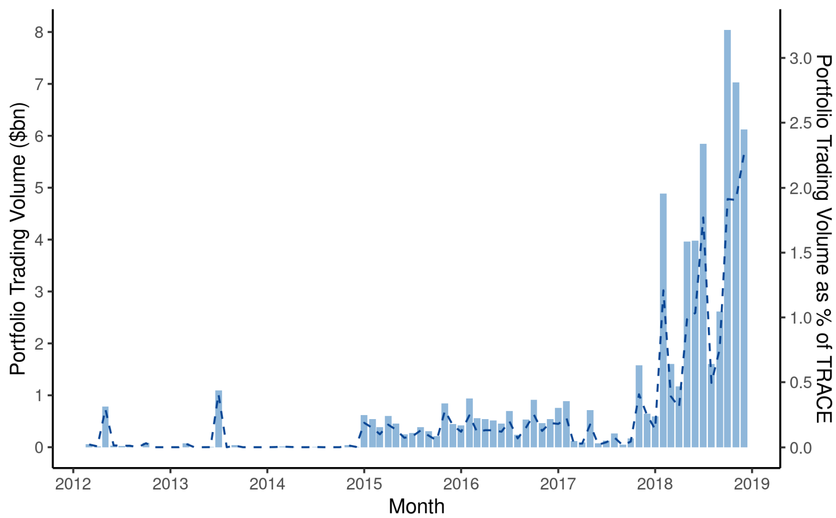

Note: This figure plots monthly corporate bond portfolio trading volume from March 2012 to December 2018. The blue bars plot the monthly portfolio trading volume (in $bn) during the period. The dashed line traces the monthly portfolio trading volume as a percentage of the total dealer-to-customer TRACE volume during the same period. I define a portfolio trade as a basket of at least 30 unique individual bonds traded at the same time by the same dealer with a customer. Data are compiled using corporate bond transaction data from the Academic TRACE dataset.

Figure 4 plots the monthly corporate bond portfolio trading volume from March 2012 to December 2018, both in absolute terms and relative to the total dealer-to-customer trading volume as reported by the Academic TRACE. Portfolio trading volume exhibits a marked increase around the end of 2017, in both absolute and relative terms.

In Figure 1 of Section 1, to show the more recent trend in corporate bond portfolio trading, I resort to the Enhanced TRACE dataset, which does not contain anonymized dealer identifiers but has more up-to-date transaction data. I obtain corporate bond transaction data between March 2012 and June 2021 from the Enhanced TRACE, and apply standard filtering and cleaning procedures as described above. Due to the lack of dealer codes, I identify portfolio trades as baskets of at least 30 unique CUSIPs traded at the same time between a dealer and a customer. In Appendix A, Figure A.5 documents the same patterns using different cutoffs for the number of unique CUSIPs used to define portfolio trades. The results for 2018 are substantially similar to those

[7] U.S. Securities and Exchange Commission Fixed Income Market Structure Advisory Committee

using the Academic TRACE. With a longer time series, the observed rise of portfolio trading is even more striking.

2.2.2 Descriptive Statistics

Table 2 presents descriptive statistics for the identified portfolio transactions both in the estimation sample and in the full sample. The focus is on the estimation sample given that the rise of corporate bond portfolio trading started around the end of 2017. The full sample statistics are presented for comparison. There are 820 portfolio trades in the estimation sample, which covers the period from January 2018 to December 2018. On average, a portfolio trade has 112 line items (unique CUSIPs) involving 80 issuers. The mean notional size of portfolio trades is $57.9 million. About one third of portfolio trades involve both dealer buying and dealer selling bonds. The majority of portfolio trades mix IG and HY securities. The average portfolio price is 99.94 cents on the dollar, with 5.08% yield-to-maturity and maturity of 10.75 years. The average transaction cost at the portfolio-level is 8.93 bps.

Table 2: Summary Statistics of Portfolio Trades

	Estimation Sample	Full Sample
Total Observations	820	1,596
Aerage Number of Line Items	112	93
Average Number of Issuers	80	71
Average Notional Size ($mn)	57.9	42.2
Mix Buy and Sell (%)	33.66	31.08
Mix IG and HY (%)	56.95	56.52
Average Portfolio Price ($)	99.94	101.34
Average YTM (%)	5.08	4.78
Average Maturity (yrs)	10.75	9.38
Average Transaction Cost (bps)	8.93	6.78

Note: This table presents the summary statistics of portfolio trades at the portfolio level. The estimation sample consists of U.S. corporate bond transaction data from the Academic TRACE dataset between January 1, 2018 and December 31, 2018. The full sample consists of U.S. corporate bond transaction data from the Academic TRACE between March 1, 2012 and December 31, 2018. A portfolio trade is defined as a basket of at least 30 unique individual bonds traded at the same time by the same dealer with a customer.

Table 3 compares corporate bond transactions executed pursuant to portfolio trading to those executed through traditional voice trading. For simplicity, I refer to the former as portfolio trades and the latter as voice trades in the table. The first two columns report the descriptive statistics for the estimation sample. For comparison, I also report in the next two columns the descriptive statistics for the full sample covering the period from March 2012 to December 2018. In the estimation sample, there are 91,857 trades involving 7,939

Table 3: Comparison of Portfolio Trades and Voice Trades

(a) Distributions by Bond and Trade Characteristics

	Estimation Sample		Full Sample	
	Portfolio Trades	Voice Trades	Portfolio Trades	Voice Trades
Total Observations	91,857	6,277,935	149,170	35,262,707
Number of CUSIPs	7,939	11,392	11,266	19,168
Rating Distribution (%)				
A and above	24.3	32.5	26.9	30.4
BBB-BB	52.6	54.8	51.7	54.2
B and below	22.8	12.3	21.0	14.8
Non-Rated	0.3	0.4	0.3	0.6
Maturity Distribution (%)				
$<$ 3 Years	11.7	27.0	13.8	20.9
3-5 Years	22.5	22.4	23.4	20.4
5-10 Years	42.6	35.6	41.5	41.4
$\geq$ 10 Years	23.2	15.0	21.3	17.3
Bond Age Distribution (%)				
$<$ 1 Year	20.0	16.1	21.9	21.5
1-3 Years	35.6	30.7	35.5	31.3
3 - 5 Years	23.2	24.8	22.7	21.8
$\geq$ 5 Years	21.2	28.3	19.9	25.5
Issue Size Distribution (%)				
$<$ \$500 mn	11.7	17.8	11.2	21.5
\$500 - 1,000 mn	39.0	30.6	37.2	31.9
$\geq$ 1,000 mn	49.3	41.6	51.6	46.5
Industry Distribution (%)				
Agriculture and Other	4.9	2.9	3.9	2.0
Energy, Mining and Utilities	13.1	12.0	13.4	16.0
Construction and Manufacturing	30.8	33.5	31.5	31.1
Trade and Transportation	10.1	12.2	10.2	10.8
Non-Financial Services	21.1	16.4	20.9	17.9
Financial Services	19.9	23.1	20.2	22.2
ETF Inclusion Distribution (%)				
Held by ETFs	96.4	91.8	94.4	76.3
Not Held by ETFs	3.6	8.2	5.6	23.7
Trade Size Distribution (%)				
Micro	26.8	61.7	27.7	60.0
Odd Lot	61.9	22.7	62.3	22.5
Round Lot	9.9	11.2	9.0	12.7
Block Trade	1.4	4.4	1.0	4.8
Trade Direction Distribution (%)				
Dealer Buy	34.0	42.3	34.4	43.2
Dealer Sell	66.0	57.7	65.6	56.8

Table 3: Comparison of Portfolio Trades and Voice Trades

(b) Sample Averages of Bond and Trade Characteristics

	Estimation Sample		Full Sample	
	Portfolio Trades	Voice Trades	Portfolio Trades	Voice Trades
Average Rating	BBB-	BBB	BBB-	BBB
Average Maturity (yrs)	9.79	7.40	9.38	8.09
Average Bond Age (yrs)	3.34	4.08	3.22	3.85
Average Issue Size ($mn)	1,169.96	1,256.33	1,200.68	1,141.41
Average Trade Size ($)	516,510	704,776	451,557	761,823

Note: This table presents descriptive statistics based on the estimation sample of corporate bond transactions in the Academic TRACE from January 1, 2018 to December 31, 2018 and the full sample from March 1, 2012 to December 31, 2018. Portfolio trades are reported transactions flagged as part of a portfolio trade defined as at least 30 unique CUSIPs transacting at the same time by the same dealer. Panel (a) reports distributions of portfolio trade line items and voice trades, by various bond-level and trade-level characteristics. All numbers in the table other than the first row are percentages. ETF constituents data are from ETF Global. Panel (b) reports sample average credit rating, maturity, bond age, issue size and trade size for portfolio trade line items and voice trades.

unique CUSIPs transacted as part of portfolio trades, compared to 6,277,935 voice trades involving 11,392 CUSIPs. Table 3 shows a notable difference between portfolio trades and voice trades in terms of trade size distributions. In the estimation sample, 61.7% of voice trades are of micro-size (below $100,000 in principal amount), while only 26.8% of portfolio trade line items are of micro-size. The average trade size of portfolio trade line items is just above $500,000. While the vast majority of portfolio trade line items are odd lots, more than 11% are round lots or block trades. The trade size distribution sets portfolio trading apart from electronic trading in the corporate bond market. For example, according to Kozora et al. (2020), the average size of an electronic trade is $82,500, including trades on 6732 ATS platforms. Thus, portfolio trades involve larger trade sizes. Moreover, at the portfolio level, the notional size of a portfolio transaction averages to $57.9 million, which is substantial.

Except for trade size distributions, portfolio trades and voice trades are substantially similar. There are a few slight differences. For example, portfolio trades tend to involve bonds with slightly lower credit ratings and longer maturities. Thus, bonds traded as part of portfolio trades seem to be slightly lower in liquidity, compared to voice-trade bonds. This is consistent with the conceptual framework in Section 4, as more effective inventory hedging allows portfolio trading to improve trading liquidity, and less liquid bonds are more easily transacted via portfolio trading while portfolio trading is less beneficial for bonds that are already very liquid. Moreover, bonds traded as part of portfolio trades are somewhat more likely to be held by ETFs. In terms of trade direction, around 34% of portfolio trade line items are dealer purchasing from investors, compared to 42% in voice trades.[8]

2.2.3 Portfolio Trading Dealers

There are 21 dealers engaging in corporate bond portfolio trading in the estimation sample. Table 4 presents the top ten dealers by portfolio trading volume. It is apparent that as of the end of 2018, corporate bond portfolio trading was highly concentrated among a few dealers. The top three dealers in portfolio trading accounted for 83.3% of total portfolio trading volume, and the top five accounted for virtually all of portfolio trading volume during 2018. In contrast, the top three voice trading dealers accounted for 26.8% of total voice trading volume and the top five accounted for 43.2% during the same period. Another observation is that with a few exceptions, the top dealers in portfolio trading also have high market share in voice trading. Thus, dealers who engage in portfolio trading are generally core dealers.

Table 4: Portfolio Trading League Table

Rank	Executing Dealer	Portfolio Trading Share	Total Market Share
1	Dealer A	58.8	10.1
2	Dealer B	12.3	8.7
3	Dealer C	12.2	0.3
4	Dealer D	9.5	0.1
5	Dealer E	4.4	8.5
6	Dealer F	1.1	7.9
7	Dealer G	0.7	8.4
8	Dealer H	0.4	0.3
9	Dealer I	0.4	7.0
10	Dealer J	0.1	0.0

Note: This table ranks the top ten dealers in corporate bond portfolio trading by volume. Portfolio trading share is a dealer's share of the total portfolio trading volume in the estimation sample. Total market share is the dealer's share of total investor-dealer trading volume in the same sample. Shares are expressed in percentages. The estimation sample consists of U.S. corporate bond transaction data from the Academic TRACE dataset between January 1, 2018 and December 31, 2018. A portfolio trade is defined as a basket of at least 30 unique individual bonds traded at the same time by the same dealer with a customer. Dealer identities are anonymized.

2.3 Proliferation of Index Derivatives and the Rise of Portfolio Trading

The rise of portfolio trading in the corporate bond market around the end of 2017 is remarkable. Around the same time, a large and active market for credit index derivatives has emerged. The post-GFC regulatory reform

[8] Dealer buy trades tend to be larger in trade size on average than dealer sell trades.

requires standardized OTC derivatives to be centrally cleared. By the end of 2012, all standardized credit index derivatives were cleared with central counterparties. As a result, a highly liquid market for standardized credit index derivatives developed over time. In Appendix A, Figure A.1 plots the average daily trading volumes for iBoxx total return swaps (TRS) for the IG index and the HY index respectively. Figure A.2 plots the open interest for LQD and HYG options. It appears that these credit index derivatives have enjoyed significant growths since 2015. For example, trading volume in iBoxx HY TRS topped $213 million per day on average in 2018, around 5% of total investor-dealer trading volume in HY bonds during the same period. As of the end of 2018, the option open interests on LQD and HYG combined were roughly 21% of the open interest on SPY, arguably the largest and mostly traded ETF.

Although many bonds have traded CDS contracts, these single-name CDS contracts are traded bilaterally and are thus often illiquid (e.g., Junge and Trolle (2013), Iercosan and Jiron (2017), Paddrik and Tompaidis (2019)). Hence, single-name CDS contracts are often ineffective hedging options for dealers managing individual bond positions. In contrast, standardized and centrally cleared credit index instruments are much more liquid. For example, Riggs et al. (2020) document that index CDS traded on swap execution facilities (SEFs) have high liquidity and low transaction costs. Moreover, since diversification at the portfolio level diminishes idiosyncratic risks, portfolios of bonds are much more easily hedged with index derivatives tracking certain risk factors.

2.4 Macro Factors

For some of my empirical tests, I use a variety of macroeconomic and financial market factors as controls. Inspired by Collin-Dufresne et al. (2007), Friewald and Nagler (2019) and He et al. (2021), I use six macro-level factors motivated by the Merton (1974) structural model. These factors are the 10-year Treasury yield (R_{10}), the squared 10-year Treasury yield (R_{10}^2), the slope of the yield curve calculated as the difference between the 10-year Treasury yield and the 2-year Treasury yield ($Slope$), the market volatility as proxied by the VIX index (VIX), the probability or magnitude of a downward jump in firm value computed from option volatility surfaces ($Jump$), and the business climate as proxied by the S&P 500 stock index return (Ret). Among these factors, Treasury yields are obtained from Gürkaynak et al. (2007) database. S&P 500 index returns are from the Center for Research in Security Prices (CRSP). VIX index values are from the CBOE. *Jump* measures are computed according to Collin-Dufresne et al. (2007) using S&P 500 index options data from Option Metrics. I further include the TED spread, obtained from St. Louis Fed, as an additional control. The TED spread is the spread between the 3-month LIBOR rate based on U.S. dollars and the 3-month Treasury bill rate, and measures the cost of funding of banks and thus proxies intermediary financial health.

2.5 Corporate Bond ETFs

I use data from the CRSP and Bloomberg to identify corporate bond ETFs traded on major U.S. exchanges. I first acquire a list of ETFs ever traded between March 2012 and June 2020 from CRSP by screening for historical share code equal to 73 (identifier for ETFs). I obtain a total of 3,087 ETFs. I then use data from Bloomberg and screen for ETFs traded on a U.S. exchange, denominated in U.S. dollars and with a fixed income asset class focus. Using information from Bloomberg, which is corroborated by security holdings data from the ETF Global and fund prospectuses, I narrow the list down to 167 ETFs with U.S. corporate bond mandates, after excluding ETFs focused primarily on Treasury securities, municipal bonds, convertible securities, preferred shares, asset-backed securities, non-U.S. and emerging markets, loans and funds of funds. Of the 167 ETFs, 61 rely on fully in-kind creation and redemption mechanism. I discuss the different types of creation and redemption mechanisms in more details in Appendix C.1.

Unlike open-ended mutual funds, ETFs are traded on exchanges throughout the trading session just like stocks. The share prices of ETFs are generally brought in line with the underlying NAVs through the in-kind creation and redemption mechanism, as described by Pan and Zeng (2019), Dannhauser and Hoseinzade (2021). Figure 5 shows the monthly total ETF arbitrage volume as well as the monthly net creation volume for the 61 corporate bond ETFs that rely solely on the in-kind creation and redemption mechanism.[9] The primary ETF creation and redemption volume steadily increased from 2014 to 2017, which is not surprising given the growth of assets in corporate bond ETFs (Figure A.3). ETF arbitrage and net creation volumes remained relatively stable during 2017 and 2018.

3 Portfolio Trading and Inventory Hedging

In this section, I establish that portfolio trading is driven by dealer inventory hedging. To address the empirical challenge that dealers' actual hedging activities are unobserved, I focus on how inventory hedging would lead to differences in dealers' intermediating behavior that are observable in the data. First, I show that while inventory

[9] I focus on the 61 ETFs with in-kind creation and redemption for two reasons. First, cash redemption and creation process is substantially similar to the subscription and redemption process of traditional open-ended mutual funds and thus have little relevance to topics studied in this paper. Second, the 61 ETFs with in-kind creation and redemption mechanism accounted for more than 80% of all assets held across the 167 corporate bond ETFs, as of December 31, 2018.

Figure 5: Corporate Bond ETF Arbitrage Volume

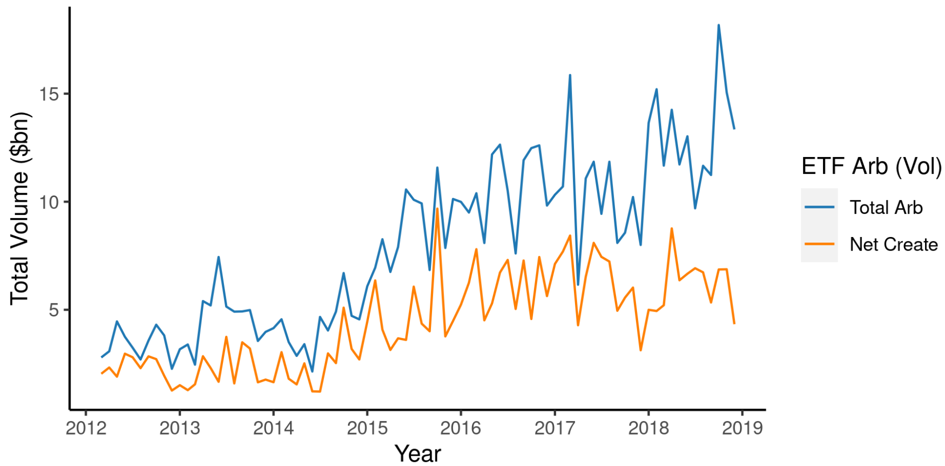

Note: This figure plots the monthly ETF arbitrage volume and net creation volume (in $ bn) of the 61 corporate bond ETFs that rely solely on the in-kind creation and redemption mechanism, for the period from March 2012 to December 2018. ETF arbitrage volume is calculated as: $|\text{Shares Outstanding}_t - \text{Shares Outstanding}_{t-1}| \times NAV_t$, or the absolute value of the product of change in shares outstanding and NAV per share. ETF net creation volume is calculated according to the formula but only for ETFs whose shares outstanding have increased. ETF shares outstanding and NAV per share data are both obtained from Bloomberg.

shocks are associated with higher transaction costs for voice-trade bonds, the transaction costs of portfolio-trade bonds do not respond to such shocks. The differential sensitivities to inventory shocks suggest that portfolio trading results in little inventory exposure, supporting the inventory hedging hypothesis. However, another possibility may be that portfolio-trade bonds are quickly transacted in offsetting trades, for example, portfolio trades are mostly pre-arranged. If this were the case, it would be observed in the data that portfolio-trade bonds turn over faster than voice-trade bonds. On the contrary, I show that bonds traded through portfolio trading in fact stay in dealer inventories much longer, thus ruling out this possibility. The slow turnover of portfolio-trade bonds is consistent with the hypothesis that portfolio trades are driven by inventory hedging, as hedging reduces dealers' incentives to offset trades quickly. Finally, I address and rule out alternative explanations.

3.1 Lack of Sensitivity to Inventory Shocks

The existing literature suggests that dealers' inventory costs are weakly convex (e.g., Amihud and Mendelson (1980)). Thus, positive inventory shocks increase dealers' inventory exposures and their marginal inventory holding costs. Intuitively, inventory shocks lead to higher inventory frictions, making it more costly for dealers to hold inventories. Higher inventory costs reduce dealers' willingness to provide liquidity, leading to higher transaction costs of voice trading. Since inventory hedging mitigates inventory frictions and reduces inventory holding costs, if portfolio trading is associated with inventory hedging, its transaction costs should exhibit low sensitivity to inventory shocks. Conversely, if the transaction costs of portfolio trading show little response to inventory shocks, then there are two likely explanations - (1) portfolio trading is associated with more effective inventory hedging, either because portfolio trades themselves are hedged or because portfolio trading results in more hedgeable inventories; or (2) portfolio trades are quickly offset in subsequent trading. In Section 3.2, I formally rule out the second explanation. Thus, to test that portfolio trading is driven by inventory hedging, I only need to show that the transaction costs of portfolio trading are less sensitive to inventory shocks than those of voice trading.

I proxy inventory shocks with issue-dealer-day level bond flows, $Flow\,Shock_{ijt}$, defined as the net purchase of bond i by dealer j on day t (in \$ millions). The bond flow shock measure quantifies the change (or shock) in dealer j's inventory of bond i on day t.[10] Due to specialization, dealer inventory management is decentralized across traders within the dealer firm (see Naik and Yadav (2003a), Nagel (2012), Cespa and Foucault (2014)). Furthermore, since corporate bonds are heterogeneous in their characteristics, it is often difficult to find a perfect substitute for any given bond. Thus, the flow shock at the issue-dealer-day level increases the marginal inventory carrying cost of the bond by the dealer, and the transaction cost in the same bond by the same dealer should increase.[11]

For robustness, I also use issue-day level bond flows as an alternative proxy for bond inventory shocks. The issue-day level bond flow measure, $Flow\,Shock_{it}$ is calculated as the net purchase of bond i by the dealer sector on day t and is the issue-dealer-day level flow shocks aggregated across all dealers. Since dealers often turn to the inter-dealer market to unload acquired positions, issue-dealer-day level flow shock measures may be subject to noise, despite their granularity. In contrast, the issue-day level inventory shocks proxy how constrained the dealer sector is as a whole with respect to a given bond. Since an aggregate shock in a bond decreases a dealer's incentive to buy or ability to sell in the inter-dealer market, it increases inventory costs associated with the bond, lowering liquidity provision and increasing transaction costs.

Furthermore, I instrument $Flow\,Shock_{ijt}$ with the dealer-day level bond flow shocks in comparable bonds matched by industry, rating, maturity and issue size categories. This instrumental variable approach follows the existing literature (e.g., Hasbrouck and Saar (2013), Comerton-Forde and Putniņš (2015), O'Hara and Zhou (2021b), Buti et al. (2022)). Bond flow shocks in comparable bonds are unlikely to directly affect the transaction

[10] Since bonds stop trading at maturity or when called, the calculation of bond-dealer-day level bond flows (as well as bond-day level bond flows) automatically adjusts for calls and maturities.

[11] See Section 4 for a model.

costs in a given bond, as flows are usually driven by the idiosyncratic liquidity needs of different investors and investor bond holdings tend to be disperse. Thus, the exclusion restriction is plausibly satisfied. On the other hand, since flow shocks in comparable bonds increases a dealer's exposure to certain common risk factors, the dealer becomes more averse to provide liquidity in the given bond. The fact that the first stage coefficient is positive and highly statistically significant, with an F-statistic of 61, further validates the relevance condition (see Stock et al. (2002)).

To test whether bonds traded as part of portfolio trades have significantly reduced sensitivity to inventory shocks, I conduct a split-sample analysis by estimating the sensitivities of transaction costs to bond flow shocks for the voice-trading sub-sample and the portfolio-trading sub-sample respectively.[12] The voice-trading sub-sample consists of investor-dealer trades not identified as part of a portfolio trade, while the portfolio-trading sub-sample includes all trades identified as part of a portfolio trade. The baseline specification is:

$$Cost_{ijst}^{k} = \beta FlowShock_{ijt} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijkst} \tag{2}$$

where i indexes bond issue, j indexes executing dealer, s indexes trade size (micro, odd lot, round lot or block), and t indexes day. The superscript k indicates the trade type sub-sample, where trade type is either portfolio trade or voice trade. The outcome variable of interest is $Cost_{ijst}^{k}$, the average transaction cost of a type-k trade with size s in bond i executed by dealer j on day t. It is calculated by averaging the bond-dealer-size-day level transaction cost estimates, computed according to Section 2.1.1. $FlowShock_{ijt}$ is dealer j's flow shock in bond i on day t, calculated according to above. α_i, α_j, α_s and α_t are bond issue, dealer, trade size and day fixed effects, respectively. The fixed effects control for differences across bond issues, dealers, trade size buckets, and time. The standard errors are double clustered at the issue-day and the dealer-day levels. The coefficient of interest is β. A positive β means that a positive flow shock in a given bond of a given dealer increases the transaction cost of the bond by the same dealer, consistent with the story that inventory shocks increase marginal inventory holding costs and result in higher transaction costs.

If bonds traded as part of portfolio trading are unrelated to inventory hedging and thus go in dealers' inventories the same way as those traded through regular voice trading, then these portfolio-trade bonds should not have differential sensitivity to inventory shocks compared to voice-trade bonds. Conversely, if portfolio-trade bonds turn out to be less (or not) sensitive to inventory shocks, it suggests that portfolio trading results in

[12]The split sample analysis allows coefficients on fixed effects to differ in sub-samples, and thus is more flexible. I estimate an interactive model in Table A.2 of Appendix A.2.

reduced inventory exposure. There are two reasons why this could be the case. First, bonds traded as part of portfolio trades are hedged or result in highly hedgeable inventories, thus decreasing inventory costs and enabling dealers to better weather inventory shocks. Second, it may take less time to arrange offsetting trades in the case of portfolio trading. There is no reason to believe ex-ante that the latter is true. I later show that portfolio-trading bonds in fact stay in dealer inventories longer, thus formally ruling out this possibility. For robustness, I also discuss and rule out other potential explanations in Section 3.3.

Table 5 reports the regression results.[13] The coefficient estimates are all positive and highly statistically significant across specifications for the voice-trading sub-sample. However, the coefficient estimates for the portfolio-trading sub-sample are all negative but statistically indistinguishable from zero. Columns (1)-(3) regress transaction cost measures on the issue-day level bond flow shocks, while controlling for various fixed effects. Column (4) presents the results for the baseline regression. Since transaction costs vary widely across trade size buckets, it is important to include trade size fixed effects. And given that the dealers who engage in portfolio trading tend to be large core dealers, and thus likely differ from other dealers on unobserved characteristics, it is key to also control for dealer-level differences. The coefficient estimate implies that a $1 million positive shock in bond flow for a given dealer is associated with a 0.58 bps increase in transaction cost for bonds transacted via voice trading. Given that the standard deviation of $Flow_{ijt}$ is $2.1 million, a back-of-envelope calculation suggests that a one standard deviation increase in bond flow shock is associated with a 1.2 bps increase in voice-trading transaction cost, or more than 5% of the mean transaction cost. In contrast, portfolio-trade bonds do not seem to respond to bond flow shocks. Column (6) instruments issue-dealer-day level flow shocks with flow shocks in comparable bonds. The 2SLS results confirm the conclusions of the baseline results from column (4).

A potential concern may be that portfolio-trade bonds are fundamentally different from voice-trade bonds, and are thus differentially sensitive to bond flow shocks. For example, portfolio-trade bonds may be higher quality or more liquid, thus incurring lower inventory carrying costs and hence the lower sensitivity to flow shocks. However, Table 3 suggests that this is not the case, as portfolio-trade bonds and voice-trade bonds are similar in bond characteristics. To further address this concern, I saturate specification (2) by replacing the issue, trade size, and day fixed effects with interacted issue-size-day fixed effects, which compare trades of the same size bucket in the same bond on the same day. The interacted fixed effects thus also control for any potential time-varying impact of bond-level and trade-level characteristics on transaction costs, which may be

[13] For robustness, I also compare portfolio-trade bonds to voice-trade bonds within the sub-sample of portfolio-trading dealers. The results are presented in Table A.3 in Appendix A.2. The results are substantially similar.

Table 5: Differential Sensitivities of Transaction Costs to Bond Flow Shocks

	Transaction Cost (bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
Voice-Trading:						
$Flow\,Shock_{it}$	0.418***	0.210***	0.053**			
	(0.100)	(0.053)	(0.020)			
$Flow\,Shock_{ijt}$				0.584***	0.241***	2.950***
				(0.066)	(0.038)	(1.052)
R^2	0.152	0.116	0.115	0.116	0.396	-0.002
Observations	1,422,753	2,330,761	4,574,528	4,574,528	3,214,116	4,574,528
Portfolio-Trading:						
$Flow\,Shock_{it}$	-0.052	-0.026	-0.014			
	(0.150)	(0.130)	(0.128)			
$Flow\,Shock_{ijt}$				-0.383	-0.399	-2.668
				(0.904)	(0.579)	(4.112)
R^2	0.230	0.227	0.230	0.230	0.617	-0.004
Observations	83,475	86,410	88,554	88,554	4,212	88,554
Issue FE	Yes	Yes	Yes	Yes	No	Yes
Day FE	Yes	Yes	Yes	Yes	No	Yes
Size FE	No	Yes	Yes	Yes	No	Yes
Dealer FE	No	No	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	No	No	No	Yes	No
Specification	FE	FE	FE	FE	FE	2SLS

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table reports the split-sample regression results examining the differential sensitivities in transaction costs of portfolio-trade bonds and voice-trade bonds to bond flow shocks. Bond flow shock is a dealer's net purchase of the bond on a given day. The estimation sample covers U.S. corporate bond transaction data reported by the Academic TRACE from January 1, 2018 to December 31, 2018. The voice-trading sub-sample consists of investor-dealer trades not identified as part of a portfolio trade, while the portfolio-trading sub-sample includes all trades identified as part of a portfolio trade. The outcome variable in column (1) is $Cost_{it}$, the average transaction cost in bond i on day t. The outcome variable in column (2) is $Cost_{its}$, the average transaction cost in bond i of trade size s on day t. The outcome variable in column (3)-(6) is $Cost_{ijts}$, the average transaction cost of a trade with size s in bond i executed by dealer j on day t. The explanatory variable $Flow\,Shock_{ijt}$ is the net purchase of bond i by dealer j on day t (in \$mn). $Flow\,Shock_{it}$ is issue-dealer-day level flow shocks aggregated across all dealers (in \$mn). Column (5) further saturates the model with issue-size-day and dealer fixed effects. Column (6) instruments $Flow\,Shock_{ijt}$ with flow shock in comparable bonds matched by industry, rating, maturity and issue size categories, in a 2SLS specification. In columns (1) and (2), standard errors are clustered at the issue and the day levels. In columns (3)-(6), standard errors are double clustered at the issue-day and the dealer-day levels.

of concern if, for example, time-varying risk aversion of dealers may lead to differential inventory sensitivities to these bond-level and trade-level characteristics over time. The coefficient estimate is reported in column (5) of Table 5. The coefficient estimate is slightly lower for the voice-trading sub-sample, but remains positive and statistically significant at the 1% level.

The above results show that positive bond flow shocks increase transaction costs of bonds traded through voice trading. However, portfolio-trade bonds are not sensitive to bond flow shocks. That is, consistent with theory, flow shocks increase inventory frictions and holding costs for dealers engaging in voice trading, resulting in higher transaction costs. In contrast, bonds traded as part of portfolio trades do not exhibit the same sensitivity to inventory frictions. Thus, the results in this section support the hypothesis that portfolio trades are associated with significantly reduced, if at all, inventory exposure. To further support that portfolio trading is associated with inventory hedging, I proceed to focus on bond turnover for portfolio trades and voice trades.

3.2 Slower Bond Turnover

Without hedging, dealers must reduce their inventory exposure by entering into offsetting trades. Since holding inventory positions is costly, dealers are incentivized to seek out counterparties with whom they can enter into inventory-reducing transactions. If, on the other hand, portfolio trades are largely hedged, then dealers' incentive to search for such counterparties is reduced.[14] Thus, if portfolio trading is tied to inventory hedging, then bonds traded as part of portfolio trades tend to turn over slower compared to bonds transacted via voice trading. Equivalently, the time it takes to turn over portfolio-trade bonds is longer relatively speaking.

By testing this hypothesis, I also address an important potential concern of my prior analysis. In Section 3.1, I show that while bond flow shocks increase the transaction costs of voice trades, portfolio trades do not seem to respond (if not negatively respond) to the same bond flow shocks, suggesting that portfolio trades are associated with significantly reduced, or little, inventory exposure. However, the reduced inventory exposure may be due to hedging or dealers' higher ability to enter into offsetting trades. Thus, the potential concern is that portfolio trades do not benefit from more effective hedging technology, instead, bonds transacted through portfolio trading spend less time in dealer inventories, resulting in lower inventory exposure and lower sensitivity to bond flow shocks. This concern can be immediately ruled out by testing that bonds traded as part of portfolio trades have longer turnover times than bonds traded through voice trading. To do this, I estimate the following

[14] This is consistent with the endogenous search story of Goldstein and Hotchkiss (2020).

baseline specification:

$$Turnover^b_{ijkst} = \beta Port\,Trade_{ijkst} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijkst} \tag{3}$$

where i indexes bond issue, j indexes dealer, k indexes trade category (either voice trade or portfolio trade), s indexes trade size, and t indexes day. The dependent variable $Turnover^b_{ijkt}$ is the average turnover time (in days) in bond i of size s purchased by dealer j in a type-k trade on day t. I first focus on bond turnover times of dealer-buy trades in my main specification for reasons explained in Section 2.1.2. However, I also present results for the full sample containing both dealer-buy and dealer-sell trades.The explanatory variable $Port\,Trade_{ijkst}$ is an indicator variable equal to one if the trade is part of a portfolio trade and zero otherwise. α_i, α_j, α_s and α_t are issue, dealer, trade size, and day fixed effects, respectively. The standard errors are double-clustered at the issue-day and the dealer day levels. The coefficient of interest is β. A positive β means that bonds traded as part of portfolio trades are associated with slower turnover compared to voice-trade bonds. The results are presented in Table 6.

Columns (1)-(3) report results for the sub-sample containing only dealer-buy trades. Column (1) estimates the baseline specification (3). The coefficient estimate is positive and statistically significant at the 1% level. It suggests that portfolio-trade bonds tend to turn over slower and stay in dealer inventories for 5.3 days longer on average. Given that the mean turnover time for dealer-buy voice trades is about 10.7 days, an estimated coefficient of 5.3 days means that the turnover in portfolio-trade bonds is about 50% slower than that in voice-trade bonds, even after controlling for differences across bond issues, dealers, trade size buckets and time. Column (2) again saturates the baseline specification by replacing the issue, trade size and day fixed effects with the interacted issue-size-day fixed effects. The coefficient estimate is slightly lower but still both highly statistically significant and economically meaningful. In column (3), I conduct a matched sample analysis by including industry-rating-maturity-amount-size-dealer-week interacted fixed effects. This effectively compares portfolio trades and voice trades of similar trade sizes in bonds with similar characteristics and executed by the same dealer during the same week. I also include various macro controls from Section 2.4 in the regression. Again, the coefficient estimate is little changed from the baseline specification.

I estimate the baseline regression for the full sample containing both dealer-buy and dealer-sell trades in column (4). Since the bond turnover measure is constructed for both dealer-buy and dealer-sell transactions, I also add the trade side fixed effects. The coefficient estimate is slightly lower than that of column (1), suggesting that the bond turnover differential between portfolio-trade bonds and voice-trade bonds is slightly smaller for

Table 6: Slower Bond Turnaround of Portfolio-Trade Bonds

	Bond Turnover (days)					
	(1)	(2)	(3)	(4)	(5)	(6)
Port Trade	5.298***	3.358***	6.016***	3.725***	3.493***	3.617***
	(0.729)	(0.419)	(0.658)	(0.494)	(0.494)	(0.481)
Issue FE	Yes	No	No	Yes	Yes	Yes
Day FE	Yes	No	No	Yes	Yes	Yes
Size FE	Yes	No	No	Yes	Yes	Yes
Dealer FE	Yes	Yes	No	Yes	Yes	Yes
Side FE	No	No	No	Yes	Yes	Yes
Issue-Size-Day FE	No	Yes	No	No	No	No
Matched	No	No	Yes	No	No	No
R^2	0.185	0.508	0.340	0.194	0.166	0.208
Observations	2,144,084	1,131,955	1,996,702	4,628,665	1,270,436	1,999,335
Sample	Dlr-Buy	Dlr-Buy	Dlr-Buy	Full	HY	PTDealer

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table reports results from regressing bond turnaround time (measured in days) on the portfolio trade dummy that equals one if the trade is part of a portfolio trade and zero otherwise. Columns (1)-(4) contain results using the sub-sample with only dealer-buy trades. Columns (5) and (6) report results using the full sample. The full estimation sample covers U.S. corporate bond transaction data reported by the Academic TRACE from January 1, 2018 to December 31, 2018. Column (1) runs a pooled regression while controlling for trade-level, issue-level and macro factors. Column (2) includes issue, size, day and dealer fixed effects. Column (3) augments the specification by including issue-size-dealer fixed effects and day fixed effects. Column (4) and (6) match similarly-sized trades in bonds of the same industry and credit rating, by the same dealer during the same week. Column (5) replicates the regression in column (2) using the full sample, with the addition of trade side fixed effects. The dependent variable in columns (1)-(4) is $Turnaround^{b}_{ijkt}$, the average turnaround time in bond i purchased by dealer j in a type-k trade on day t. The dependent variable in columns (5) and (6) is the average turnaround time in bond i traded by dealer j in either direction, in a type-k trade on day t. Standard errors are double clustered at the issue-day and the dealer-day levels.

dealer-sell trades. The coefficient estimate for the full sample is still highly significant both statistically and economically. In columns (5) and (6), I repeat the same analysis for the sub-sample of high-yield bonds and the sub-sample containing only portfolio-trading dealers, respectively. Results are little changed.

In summary, the results in Sections 3.1 and 3.2 show that while inventory shocks increase transaction costs in voice trading, portfolio-trade bonds show little sensitivity to inventory shocks, suggesting significantly reduced inventory exposure of portfolio-trade bonds. However, this reduced inventory exposure is not due to dealers' increased ability to enter into offsetting transactions. On the contrary, portfolio-trade bonds tend to stay in dealer inventories longer compared to voice-trade bonds. These results combined support the hypothesis that portfolio trading is tied to more effective inventory hedging at the portfolio level.

3.3 Ruling Out Alternative Explanations

3.3.1 Alternative Explanation 1: ETF Creation and Redemption

One alternative possibility is that instead of hedging, dealers engage in portfolio trading to leverage the ETF creation and redemption processes and transform liquidity. Some recent studies (e.g., Shim and Todorov (2021) and Reilly (2021)) further suggest that authorized participants (APs), who are also bond dealers, tend to have some leeway in negotiating the final creation or redemption baskets with the ETF sponsors. However, this mechanism likely plays an insignificant role in portfolio trading during my estimation sample period, as after 2006 and before the implementation of Rule 6c-11 in December 2019, the SEC imposed tight restrictions on creation and redemption baskets, allowing limited room for non-pro rata baskets (or custom baskets).[15]

To rule out this possibility, I follow Shim and Todorov (2021) and Reilly (2021) to construct bonds that are delivered or redeemed as part of the ETF creation and redemption process on a daily basis. I then compare these bonds to the bonds traded as part of portfolio trades. In particular, I compute the percentage of portfolio-trade bonds purchased by dealers t that are subsequently included in an ETF creation basket on $t+k$ with $k = 0, 1, ..., 30$. I then repeat the same calculation for dealer-sell bonds in portfolio trades. I find that the overlap is 20-25% for portfolio-trade bonds purchased by dealers, and 25-30% for portfolio-trade bonds sold by dealers. In other words, 70-80% of bonds in portfolio trades are not delivered or redeemed through the ETF creation and redemption process. Given that on average 373 bonds are traded as part of portfolio trading and 1,314 bonds are included in ETF baskets each day, an overlap of 20-30% is remarkably small, suggesting that the ETF creation and redemption unlikely plays any significant role in portfolio trading.

[15]For example, see A little ETF rule change that could make a big difference.

Figure 6: Overlap between Portfolio-Trade Bonds and ETF Baskets

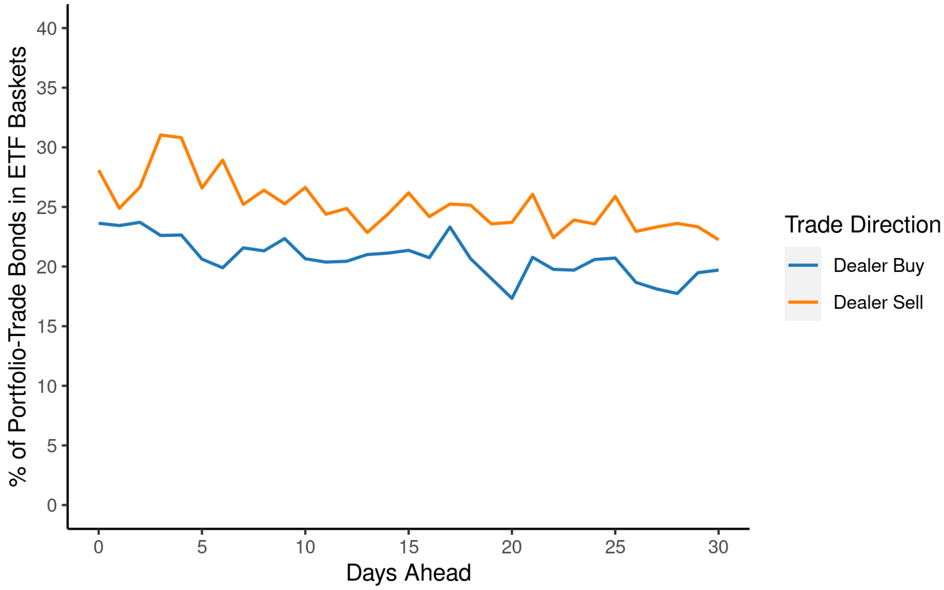

Note: This figure plots the overlap between portfolio-trade bonds and realized ETF baskets of the corporate bond ETFs with in-kind creation and redemption mechanism. Realized ETF baskets are inferred according to Shim and Todorov (2021) and Reilly (2021). The overlap for portfolio-trade bonds purchased by dealers is calculated as the percentage of these bonds that are subsequently included in an ETF creation basket with a lag of k days ($k = 0, 1, ..., 30$). Similarly, the overlap for portfolio-trade bonds sold by dealers is calculated as the percentage of these bonds that are subsequently included in an ETF redemption basket with a lag of k days. ETF holdings data are obtained from ETF Global while ETF shares data are from Bloomberg. The sample period is January 1, 2018 to December 31, 2018.

To provide further evidence, I regress daily portfolio trading volume on daily ETF arbitrage volume. I then regress daily portfolio trading volume involving dealer-buy trades on daily ETF net creation volume. The results are presented in columns (1) and (2) of Table A.4 in Appendix A.2. The coefficient estimates are statistically insignificant and R^2 are close to zero in both regressions, suggesting that there is little correlation between ETF arbitrage or net creation and portfolio trading. I also test whether transaction costs of portfolio-trade bonds and voice-trade bonds are differentially sensitive to changes in the ETF-NAV bases. Given the segmented nature between IG and HY,[16] I separately test whether IG portfolio-trade and voice-trade bonds are differentially sensitive to changes in the LQD ETF-NAV basis, and whether HY portfolio-trade and voice-trade bonds are differentially sensitive to changes in the HYG ETF-NAV basis. Since LQD and HYG are respectively the largest IG and HY ETFs by assets-under-management, their ETF-NAV bases serve as good proxies for ETF arbitrage (i.e., creation and redemption) activity in IG and HY markets respectively. A large basis suggests a lack of arbitrage capital from dealers, and low level of activity in ETF creations and redemptions. Hence, if dealers engaging in portfolio trading rely on ETF creations and redemptions to transform liquidity, they may find it

[16]For example, Ellul et al. (2011), Chernenko and Sunderam (2012) among others.

more difficult to do so when the ETF-NAV basis is large, leading to higher relative transaction costs. The results are also reported in Table A.4. The transaction costs of portfolio-trade bonds are not shown to differentially respond to changes in ETF-NAV basis.

The above results suggest that the ETF creation and redemption mechanism does not play a significant role in portfolio trading, at least during the estimation sample period. Because portfolio trading started to rapidly gain popularity around the end of 2017, the fact that the ETF creation and redemption mechanism played an insignificant role means that the rise of portfolio trading was not due to this mechanism, but instead facilitated by other considerations, namely inventory hedging.

3.3.2 Alternative Explanation 2: Private Information and Adverse Selection

Another potential explanation is that the reduced information sensitivity exposes dealers to less adverse selection, leading to improved liquidity and lower sensitivity to inventory shocks (e.g., Glosten and Milgrom (1985)). However, if the observed lower sensitivity of portfolio-trade bonds to inventory shocks were mainly due to reduced adverse selection problem, then these bonds should turn over faster in dealer inventories due to increased liquidity. The fact that portfolio-trade bonds are associated with slower bond turnover suggests that inventory hedging is the more likely explanation.

To test this empirically, I proxy private information using the 14-day realized equity volatilities ($Volatility_{it}$).[17] Drechsler et al. (2021) suggest that volatility is induced by private information, and high realized volatility reflects more private information coming to light and thus dealers become more sensitive to order flows. Hotchkiss and Ronen (2015) show that stocks and bonds contain similar firm-specific information content. On the other hand, realized equity volatilities are unlikely to affect transaction costs of corporate bonds through channels other than adverse selection. Thus, if portfolio-trade bonds' lower sensitivity to inventory shocks is due to reduced adverse selection, the differential sensitivity should be more pronounced when the realized equity volatilities are elevated. I estimate the following baseline specification:

$$Cost_{ijkst} = \beta_1 Port\,Trade_{ijkst} + \beta_2 Volatility_{it} + \beta_3 Port\,Trade_{ijkst} \times Volatility_{it} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijkst}$$

where i indexes bond issue, j indexes dealer, k indexes trade type (portfolio trade or voice trade), s indexes trade size, and t indexes day. The outcome variable is the average issue-dealer-size-day level transaction cost of a type-k trade. α_i, α_j, α_s and α_t are issue, dealer, trade size, and day fixed effects. I also saturate the specification

[17] The 14-day calculation window for realized volatilities roughly matches the average turnover of corporate bonds in Section 2.1.2.

by replacing α_i, α_s and α_t with α_{ist}, the interacted issue-size-day fixed effects. The coefficient of interest is β_3. If adverse selection is the main channel through which portfolio trading reduces inventory sensitivity, then transaction cost effect of portfolio trading should be more pronounced when there is a greater amount of private information and thus more severe adverse selection problem. That is, β_3 should be negative. The regression results are reported in Table 7. In columns (1) and (2), the explanatory variable $Volatility_{it}$ is the calculated 14-day realized equity volatility (in percentage) of issue i on day t. In columns (2) and (3), the explanatory variable $Volatility_{it}^d$ is a dummy variable created from $Volatility_{it}$ and equals one if $Volatility_{it}$ is above its sample median, and zero otherwise. None of the β_3 estimates is statistically distinguishable from zero. Hence, the transaction cost advantage of portfolio trading does not seem to be related to the amount of private information or adverse selection. Thus, the private information and adverse selection channel is unlikely the driving force.

Table 7: Differential Transaction Costs and Realized Equity Volatilities

	Transaction Cost (bps)			
	(1)	(2)	(3)	(4)
$Port\,Trade_{ijkst} \times Volatility_{it}$	0.005	-0.018		
	(0.039)	(0.034)		
$Port\,Trade_{ijkst} \times Volatility_{it}^d$			-0.223	-0.026
			(0.600)	(0.574)
Issue FE	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
Size FE	Yes	Yes	Yes	Yes
Dealer FE	No	Yes	No	Yes
R^2	0.149	0.191	0.149	0.191
Observations	2,651,629	2,651,575	2,651,629	2,651,575

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table reports the differential impact of realized equity volatilities on bond transaction costs for portfolio-trade bonds and voice-trade bonds. $Port\,Trade_{ijkst}$ is the portfolio trade dummy. $Volatility_{it}$ is the 14-day realized equity volatility (in percentage) of bond issue i on day t. $Volatility_{it}^d$ is an indicator variable equal to one if $Volatility_{it}$ is above its sample median, and zero otherwise. The outcome variable in columns (1) and (3) is $Cost_{ist}$, the average issue-size-day level transaction cost, while the outcome variable in columns (2) and (4) is $Cost_{ijst}$, the average issue-dealer-size-day level transaction cost. In columns (1) and (3), standard errors are two-way clustered at the issue and the day levels. In columns (2) and (4), standard errors are double clustered at the issue-day and the dealer-day levels. The estimation sample consists of corporate bond transactions from January 1, 2018 to December 31, 2018.

3.3.3 Alternative Explanation 3: Growth of Index Funds

Furthermore, the rise of portfolio trading could be related to the increase in assets-under-management of index funds, resulting in increased demand for trading multiple securities at the same time. Although an interesting demand-side story, the explanation cannot directly explain why portfolio-trade bonds and voice-trade bonds are differentially sensitive to inventory shocks, and why portfolio-trade bonds are associated with slower bond turnover. Thus, it is important to understand the supply-side (i.e. dealer-side) reasons. The focus of this paper is to provide such supply-side explanations for why dealers choose to engage in portfolio trading. The analyses above point to inventory hedging as the main reason.

There is, however, an indirect effect of the growth in index funds. As index and passive funds are more concerned with broad market (or factor) exposures instead of idiosyncratic risks, credit index products may be particularly useful to these funds for gaining exposures to a certain index. For example, index funds with recent inflows can quickly invest in total return swaps (TRS) or ETF options, with minimal capital tie-up. Investing in these index derivatives allows the index funds to minimize tracking error while maintaining the ability to invest their subscription proceeds in corporate bonds over time. As a result, index funds are natural counterparties to dealers in inventory hedging transactions. The growth of index funds in the corporate bond space thus has likely facilitated the development and trading of credit index derivatives, making inventory hedging more accessible for dealers at the portfolio level. While this itself is an interesting question to pursue, this paper takes the proliferation of credit index derivatives and the increasing availability of portfolio-level hedging instruments as exogenous, and instead focuses on how inventory hedging affects liquidity provision and transaction costs.

4 Conceptual Framework

Section 3 provides empirical evidence that corporate bond portfolio trading is driven by dealer inventory hedging. In this section, I develop a conceptual framework for dealer inventory hedging and study how it affects dealer liquidity provision and transaction costs in a stylized dynamic search model. The model is built upon Duffie et al. (2005). I extend their model in two dimensions. First, the inter-dealer market is bilateral and frictional. The search and bargaining frictions in the inter-dealer market leads to dealer inventory imbalances, which are costly. Second, there is a hedging technology that reduces inventory holding costs. However, dealers are subject to i.i.d. shocks to their access to this hedging technology. The hedging access shocks generate inter-dealer trading in the steady state. Thus, not only does hedging affect liquidity and transaction costs, the effect depends on inter-dealer search frictions and inventory frictions.

4.1 The Setup

Time is continuous and infinite, with $t \in [0,\infty)$. There are two types of risk-neutral agents, a unit measure of investors and a unit measure of dealers, all of whom are infinitely-lived. There is no entry or exit of investors and dealers. All agents discount time at rate $r > 0$. Consider one asset (a tree), which produces a unit flow of a perishable good called fruit. There is also a numeraire general consumption good in this economy, which is produced and consumed by all agents. The tree is durable and indivisible, and has a fixed supply of $s \in \left(0, \frac{1}{2}\right)$. That is, the supply of the tree accounts for a relatively small fraction of total available capital. The tree is traded among the agents in a decentralized market intermediated by the dealers. There is no trading market for the fruit, and the only way to consume the fruit is through the ownership of a tree.

Investors. Initially, a measure s of investors are endowed with one unit of the tree, and the dealers do not hold any tree at time 0. Investors can hold either 0 or 1 unit of the tree and cannot short sell. When holding one unit of the tree, an investor derives a utility flow ε from fruit consumption. The instantaneous utility function of an investor has a quasilinear form $q\varepsilon + C$ where $q \in \{0,1\}$ is the holding of the investor and C is the investor's consumption of the numeraire good. There are two types of investors based on their utility flow from consuming the fruit, "high" or "low" types. A high-type investor derives a flow utility ε_h, while a low-type investor derives a flow utility ε_l ($\varepsilon_h > \varepsilon_l$). An investor's type switches from low to high with intensity $\frac{1}{2}$, and from high to low with intensity $\frac{1}{2}$.

The time-varying investor types, as are standard in the literature, generate trading in the steady state. There are different interpretations of the investor types. One interpretation reflects investors' heterogeneous liquidity needs. For example, low-type investors require liquidity to meet fund redemptions or invest in a more attractive alternative opportunity, while high-type investors want to invest the proceeds from recent subscriptions or simply lack alternative options. The time-varying aspect thus represents investors' liquidity shocks, as an investor receiving subscriptions and looking to invest at one point in time may be faced with redemptions at another point. Another interpretation is that investors have heterogeneous valuations, and high-type investors value the asset more than the low-type investors. This interpretation may be further micro-founded by investors' heterogeneous beliefs about the distribution of future payoffs of the asset.

Dealers. Dealers can also hold either 0 or 1 unit of the tree on their balance sheets. Dealers holding one unit of the tree are long one unit. I call these dealers "long" dealers. Dealers who hold zero unit of the tree have no positions and are thus called "nil" dealers. Dealers do not derive utility from the fruit. Hence, dealers' role in the market is to facilitate trades, rather than act as investors and derive utility directly from the tree and the flow of fruit it produces. This assumption seems plausible, as the Volcker Rule outright prohibits bank dealers from engaging in proprietary trading post-GFC, while other regulatory reforms have significantly constrained dealers' balance sheet space and their ability to participate in profitable investing and arbitrage opportunities (e.g., Du et al. (2017)).

Hedging Technology. Without hedging, dealers incur inventory holding cost $c > 0$, which represents the disutility from inventory imbalances. Carrying unhedged inventories is costly due to VaR constraint and career risk of the dealers' traders, as well as costly monitoring and compliance. A high c means that dealers face more inventory frictions. There is a hedging technology, which decreases dealer inventory cost. A dealer either has access to the hedging technology or does not have access to it. Upon meeting with an investor, the dealer receives an i.i.d. access "shock", and the ex-ante probability that the dealer has access to hedging is π.[18] Dealers learn about whether they have access to the hedging technology with a delay. Therefore, dealers are ex-ante identical in the sense that they do not know whether they have access to hedging when contacted by investors, but their access status is revealed after they post prices and the investors disclose their trading intentions. Since there is a continuum of investors and a continuum of dealers, the fraction of dealers who have access to the hedging technology is π. Thus, π represents the accessibility of the hedging technology, and a higher π means that it is easier to hedge.

Hedging lowers inventory holding cost. For simplicity, I assume that a hedged position does not incur inventory holding cost. This assumption is without loss of generality. In practice, hedged positions generally incur some holding costs, as hedging is often imperfect, and dealers face other firm-level constraints, such as balance sheet space constraints, funding costs, and regulatory capital requirements. However, in the model, these imperfections can be subsumed by the stochastic accessibility of the hedging technology, π. While π is modeled as the accessibility of the hedging technology, it can also be interpreted as hedging cost or effectiveness of hedging. Lower cost (higher effectiveness) of hedging means that dealers are more likely to find it economical

[18] I assume that hedging access shock occurs only when dealers are met with investors, and not when dealers meet each other. It is assumed that dealers who meet with investors but ultimately do no trade also have their hedging access re-drawn. The access shock may be interpreted as new information regarding the asset value and thus hedging availability, which can be inferred from investor flows.

to hedge, resulting in a higher π, while higher cost (lower effectiveness) of hedging is associated with lower usage of hedging and thus a lower π.

Trading. Trading occurs in a decentralized market intermediated by dealers. That is, investors must trade with dealers, and cannot trade among themselves, while dealers can trade with both investors and other dealers. In my model, I focus on principal trading. That is, dealers "make a market" and trade with investors by utilizing their balance sheets. This is in contrast to agency trading, where a broker-dealer simply looks for a counterparty on behalf of the investor and only executes a trade when two sides are matched.

Figure 7: Illustration of Timing

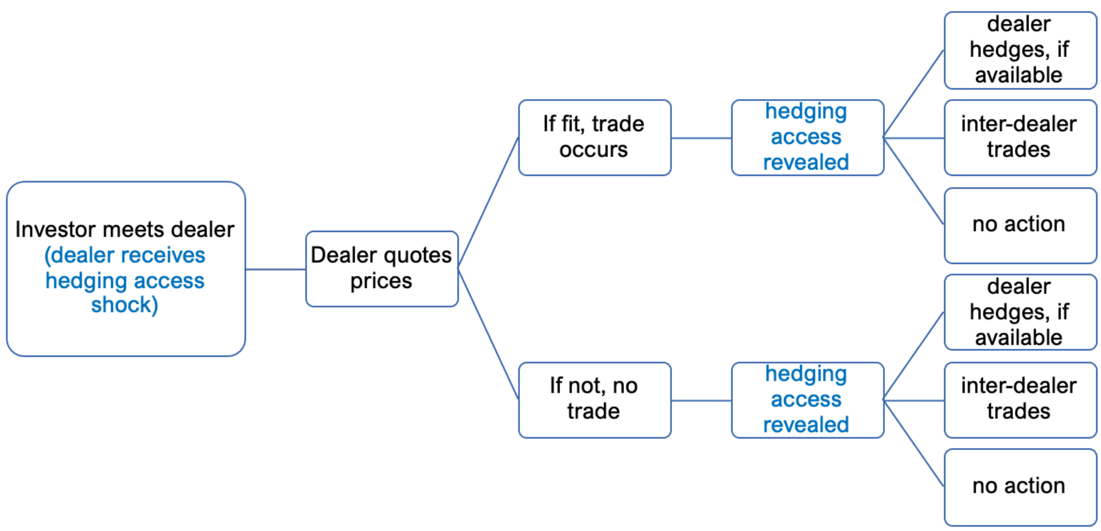

Note: This figure illustrates the timing of trading activities and information revelation about hedging access.

The sequence of trading activities and timing work as follows (also see Figure 7). Investors with trading needs search for dealers to trade. These investors randomly contact dealers and are met with a dealer at independent Poisson arrival times with intensity λ. This captures the bilateral and sequential nature of OTC trading. Dealers who are contacted by investors receive hedging access shocks, with their access to the hedging technology re-drawn from i.i.d. Bernoulli distribution with parameter π. Upon meeting with investors, these dealers must commit to bid and ask prices which are determined through Nash bargaining, with dealers' bargaining power being θ. It is then revealed whether they have access to the hedging technology.[19] Due to the holding restrictions on investors and dealers, a buying investor can trade only with a long dealer, and a

[19]This ensures that dealers are ex-ante identical and the bid and ask prices do not depend on the outcome of dealers' hedging access. However, all dealers who are contacted by investors experience the hedging access shock, even though some of them do not end up transacting.

selling investor can trade only with a nil dealer. I assume that there is always gain from trade when a buying investor meets a long dealer and when a selling investor meets a nil dealer. This assumption rules out the uninteresting case that investors never trade. Dealers may then hedge if the hedging technology is accessible to them, or engage in inter-dealer trading to unload their acquired positions. The inter-dealer market is frictional. A dealer contacts other dealers at random and is met with a dealer at independent Poisson arrival times with intensity ρ.[20] The inter-dealer market price is also determined through Nash bargaining, and dealers have equal bargaining powers.[21]

4.2 Trading Dynamics and Equilibrium

Let σ denote the investor type, with $\sigma \in \{ho, hn, lo, ln\}$, where the letters “h" and “l" represent high-type investor and low-type investor, and the letters “o" and “n" correspond to owning and not owning the tree respectively. There is a continuum of investors and let $\mu_\sigma(t)$ denote the fraction of the investors that are of type σ at time t. Note that because there is a unit measure of investors, μ_σ is equivalent to the measure of investors of type σ. Thus,

$$\mu_{ho}(t) + \mu_{hn}(t) + \mu_{lo}(t) + \mu_{ln}(t) = 1 \tag{4}$$

Let ξ denote the dealer type, then $\xi \in \{l, n\}$, representing long dealer and nil dealer respectively. There is a continuum of dealers and let $\nu_\xi(t)$ denote the fraction (and equivalently, measure) of the dealers that are of type ξ at time t. Thus,

$$\nu_l(t) + \nu_n(t) = 1 \tag{5}$$

Market clearing requires

$$\mu_{ho}(t) + \mu_{lo}(t) + \nu_l(t) = s \tag{6}$$

[20]Empirically, the search friction in the inter-dealer market is generally lower than the search friction in the investor-dealer trading market, that is, $\rho > \lambda$.

[21]The model makes multiple simplifying assumptions, in order to focus on the key mechanism of inventory hedging. The model's main results are robust to a more general environment without such assumptions.

4.2.1 Investor-Dealer Trading

At time t, investors contact dealers to trade. In equilibrium, low-type owner investors wish to sell and high-type non-owner investors wish to buy, while low-type non-owners and high-type owners do not trade. Due to the holding restrictions, low-type owners can only sell to dealers who do not own the tree, while high-type non-owners can only purchase from dealers who already own a unit of the tree. The change in the fraction μ_σ for each investor type σ depends on both trading and the transition of types. For example, consider a low-type investor owning one unit of the tree (type lo) at time t. The investor contacts dealers to trade and transitions to a low-type non-owner if the trade is successfully executed. Since only a nil dealer is able to purchase a unit of the tree, the type lo investor meets a willing dealer at an intensity $\lambda \nu_n(t)$, and the reduction of the measure of type lo investors due to trading is thus $\lambda \mu_{lo}(t)\nu_n(t)$. On the other hand, $\frac{1}{2}\mu_{lo}(t)$ of type lo investors transition into type ho investors and $\frac{1}{2}\mu_{ho}(t)$ of type ho investors transition to become type lo investors. The same logic applies to all four investor types. Thus, the changes in investor type measures are given by:

$$\begin{aligned}
\dot{\mu}_{lo}(t) &= -\lambda \mu_{lo}(t)\nu_n(t) - \frac{1}{2}\mu_{lo}(t) + \frac{1}{2}\mu_{ho}(t) \\
\dot{\mu}_{hn}(t) &= -\lambda \mu_{hn}(t)\nu_l(t) - \frac{1}{2}\mu_{hn}(t) + \frac{1}{2}\mu_{ln}(t) \\
\dot{\mu}_{ho}(t) &= \lambda \mu_{hn}(t)\nu_l(t) - \frac{1}{2}\mu_{ho}(t) + \frac{1}{2}\mu_{lo}(t) \\
\dot{\mu}_{ln}(t) &= \lambda \mu_{lo}(t)\nu_n(t) - \frac{1}{2}\mu_{ln}(t) + \frac{1}{2}\mu_{hn}(t)
\end{aligned} \tag{7}$$

4.2.2 Inter-Dealer Trading

In equilibrium, long dealers who cannot hedge their positions wish to sell their positions to nil dealers who are able to hedge, since there is gain from trade. If the trade takes place, the units of the tree simply change hands among dealers and the measure of long dealers remains unchanged. Thus, the change in the fraction ν_ξ for each dealer type ξ depends only on trading with investors. Consider long dealers at time t. The reduction of the measure of long dealers due to trading with investors is $\lambda \mu_{hn}(t)\nu_l(t)$, since investor initiate trades during investor-dealer trading and a long dealer is met with a buying investor (type hn) with intensity λ. A nil dealer becomes a long dealer if it purchases a unit of the tree from a selling investor (type lo). The measure of long dealers increases by $\lambda \mu_{lo}(t)\nu_n(t)$ from nil dealers trading with selling investors. The changes in the measures of different types of dealers are thus:

$$\begin{aligned}\dot{\nu}_l(t) &= -\lambda\mu_{hn}(t)\nu_l(t)+\lambda\mu_{lo}(t)\nu_n(t)\\ \dot{\nu}_n(t) &= -\lambda\mu_{lo}(t)\nu_n(t)+\lambda\mu_{hn}(t)\nu_l(t)\end{aligned} \tag{8}$$

4.2.3 Steady State

I focus on the steady state equilibrium. I first show the existence, uniqueness and stability of the steady state. I then present several testable predictions of this model.

Proposition 1. ***(Steady State Equilibrium)*** *There exists a unique steady state solution* $(\bar{\mu}_{lo},\bar{\mu}_{hn},\bar{\mu}_{ln},\bar{\mu}_{ho},\bar{\nu}_l,\bar{\nu}_n)$ *to* (4)-(8)*, such that the measure of each type of agents remains constant over time. The steady-state equilibrium is (locally) asymptotically stable.*

Proof. See Appendix B.1 □

Since the system converges to a unique steady state that is locally asymptotically stable, in the following equilibrium analysis, agent measures can be taken as given. For notional ease, I omit the upper-bars of steady state measures for investors and dealers. Let $V_{\sigma(t)}(t)$ be the value function representing the maximum utility of an investor of type $\sigma(t)\in\{ho,hn,lo,ln\}$ at time t. The Hamilton-Jacobi-Bellman (HJB) equations for the value functions (suppressing the time argument t) are:

$$\begin{aligned}\dot{V}_{lo} &= rV_{lo}-\frac{1}{2}(V_{ho}-V_{lo})-\lambda\nu_n(B+V_{ln}-V_{lo})-\varepsilon_l\\ \dot{V}_{hn} &= rV_{hn}-\frac{1}{2}(V_{ln}-V_{hn})-\lambda\nu_l(V_{ho}-V_{hn}-A)\\ \dot{V}_{ho} &= rV_{ho}-\frac{1}{2}(V_{lo}-V_{ho})-\varepsilon_h\\ \dot{V}_{ln} &= rV_{ln}-\frac{1}{2}(V_{hn}-V_{ln})\end{aligned}$$

where B and A are bid and ask prices quoted by dealers, and are determined through Nash bargaining. Consider a high-type non-owner investor purchasing one unit of the tree from a long dealer. If the trade occurs, the payoff to the investor is $V_{ho}-A$. If the trade does not occur, the investor continues to search and generates payoff of V_{hn}. Thus, the net payoff to the investor from trading is $V_{ho}-A-V_{hn}$. In contrast, if the trade occurs, the dealer receives the ask price A. Let M_l denote the long dealer's reservation value, i.e., the net utility the dealer can expect to generate if the trade does not take place. Then, the dealer's net payoff from trading is $A-M_l$. With

$\theta \in (0,1)$ being the dealer's bargaining power against the investor, the ask price is determined as:

$$A = \arg\max_A (V_{ho} - A - V_{hn})^{1-\theta}(A - M_l)^{\theta} = \theta(V_{ho} - V_{hn}) + (1-\theta)M_l \tag{9}$$

Similarly, let M_n denote the reservation value of the nil dealer,

$$B = \arg\max_B (B + V_{ln} - V_{lo})^{1-\theta}(M_n - B)^{\theta} = \theta(V_{lo} - V_{ln}) + (1-\theta)M_n \tag{10}$$

In the steady state, $\dot{V}_\sigma = 0$ for any investor type $\sigma \in \{ho, hn, lo, ln\}$. Hence,

$$\begin{aligned}
0 &= rV_{lo} - \frac{1}{2}(V_{ho} - V_{lo}) - \lambda \nu_n (B + V_{ln} - V_{lo}) - \varepsilon_l \\
0 &= rV_{hn} - \frac{1}{2}(V_{ln} - V_{hn}) - \lambda \nu_l (V_{ho} - V_{hn} - A) \\
0 &= rV_{ho} - \frac{1}{2}(V_{lo} - V_{ho}) - \varepsilon_h \\
0 &= rV_{ln} - \frac{1}{2}(V_{hn} - V_{ln})
\end{aligned} \tag{11}$$

Now consider dealers' maximum achievable expected utility, or value functions, immediately after meeting with investors. Let $W_\xi^h(t)$ denote the time-t value function of a dealer of type $\xi \in \{l, n\}$ with hedging access $h \in \{c_L, c_H\}$. c_L means that the dealer is hedged so that it incurs a low inventory holding cost, which is normalized to zero. c_H means that the dealer is unhedged so that it incurs a high inventory holding cost $c > 0$. For instance, for a long dealer with access to hedging, since there is no gain from trade for this dealer in the inter-dealer market, it does not trade with other dealers in equilibrium. The dealer is contacted by an investor at rate $\lambda(\mu_{hn} + \mu_{lo})$, and upon meeting an investor, the dealer's hedging access is re-drawn. The dealer is met with a buying investor at rate $\lambda\mu_{hn}$, in which case trading takes place. Thus, the HJB equation for the long dealer with hedging is:

$$\dot{W}_l^{c_L} = rW_l^{c_L} - \lambda\mu_{hn}(A + W_n - W_l) - \lambda(\mu_{hn} + \mu_{lo})(W_l - W_l^{c_L})$$

where $W_\xi \equiv \pi W_\xi^{c_L} + (1-\pi)W_\xi^{c_H}$ denotes the ex-ante value function of a type-ξ dealer. For a long dealer without hedging, on the other hand, there is gain from trade when met with a nil dealer with hedging, and the total intensity of an inter-dealer trade occurring is $2\rho\nu_n\pi$. Thus,

$$\dot{W}_l^{c_H} = rW_l^{c_H} - \lambda\mu_{hn}(A + W_n - W_l) - \lambda(\mu_{hn} + \mu_{lo})(W_l - W_l^{c_H}) - 2\rho\nu_n\pi(P + W_n^{c_H} - W_l^{c_H}) + c$$

where P is the inter-dealer market price of the tree. Similarly, for nil dealers with and without hedging:

$$\begin{aligned}
\dot{W}_n^{c_L} &= rW_n^{c_L} - \lambda\mu_{lo}(W_l - W_n - B) - \lambda(\mu_{hn} + \mu_{lo})(W_n - W_n^{c_L}) - 2\rho\nu_l(1-\pi)(W_l^{c_L} - W_n^{c_L} - P) \\
\dot{W}_n^{c_H} &= rW_n^{c_H} - \lambda\mu_{lo}(W_l - W_n - B) - \lambda(\mu_{hn} + \mu_{lo})(W_n - W_n^{c_H})
\end{aligned}$$

Thus,

$$\begin{aligned}
\dot{W}_l &= \pi\dot{W}_l^{c_L} + (1-\pi)\dot{W}_l^{c_H} = rW_l - \lambda\mu_{hn}(A - \Delta W) - 2\rho\nu_n\pi(1-\pi)(P - \Delta W^{c_H}) + (1-\pi)c \\
\dot{W}_n &= \pi\dot{W}_n^{c_L} + (1-\pi)\dot{W}_n^{c_H} = rW_n - \lambda\mu_{lo}(\Delta W - B) - 2\rho\nu_l\pi(1-\pi)(\Delta W^{c_L} - P)
\end{aligned}$$

where $\Delta W \equiv W_l - W_n$ and $\Delta W^h \equiv W_l^h - W_n^h$ for each $h \in \{c_L, c_H\}$ are reservation values. Since access to the hedging technology is revealed after meeting with investors, dealers are ex-ante identical. It is immediate that:

$$M_l = M_n = \Delta W \tag{12}$$

which is dealers' ex-ante expected reservation value. The inter-dealer market price is determined through Nash bargaining between buying and selling dealers with symmetric bargaining powers. In equilibrium, inter-dealer trading only takes place between a long dealer without hedging and thus a high inventory cost and a nil dealer with hedging and thus a low inventory cost. The inter-dealer market price is:

$$P = \frac{1}{2}\Delta W^{c_H} + \frac{1}{2}\Delta W^{c_L} \tag{13}$$

Thus, in the steady state,

$$\begin{aligned}
rW_l^{c_L} &= \lambda\mu_{hn}(A - \Delta W) + \lambda(\mu_{hn} + \mu_{lo})(W_l - W_l^{c_L}) \\
rW_l^{c_H} &= \lambda\mu_{hn}(A - \Delta W) + \lambda(\mu_{hn} + \mu_{lo})(W_l - W_l^{c_H}) + \rho\nu_n\pi(\Delta W^{c_L} - \Delta W^{c_H}) - c \\
rW_n^{c_L} &= \lambda\mu_{lo}(\Delta W - B) + \lambda(\mu_{hn} + \mu_{lo})(W_n - W_n^{c_L}) + \rho\nu_l(1-\pi)(\Delta W^{c_L} - \Delta W^{c_H}) \\
rW_n^{c_H} &= \lambda\mu_{lo}(\Delta W - B) + \lambda(\mu_{hn} + \mu_{lo})(W_n - W_n^{c_H})
\end{aligned} \tag{14}$$

and

$$
\begin{aligned}
rW_l &= \lambda\mu_{hn}(A-\Delta W)+\rho\nu_n\pi(1-\pi)(\Delta W^{cL}-\Delta W^{cH})-(1-\pi)c \\
rW_n &= \lambda\mu_{lo}(\Delta W-B)+\rho\nu_l\pi(1-\pi)(\Delta W^{cL}-\Delta W^{cH})
\end{aligned}
\tag{15}
$$

Let $\Delta V_h \equiv V_{ho}-V_{hn}$ and $\Delta V_l \equiv V_{lo}-V_{ln}$ be the reservation values of high-type and low-type investors respectively, the system of equations from (9) to (11) can be rewritten in terms of reservations values, thus

$$
\begin{aligned}
r\Delta V_l &= \frac{1}{2}(\Delta V_h-\Delta V_l)+\lambda\nu_n(B-\Delta V_l)+\varepsilon_l \\
r\Delta V_h &= \frac{1}{2}(\Delta V_l-\Delta V_h)-\lambda\nu_l(\Delta V_h-A)+\varepsilon_h \\
A &= \theta\Delta V_h+(1-\theta)\Delta W \\
B &= \theta\Delta V_l+(1-\theta)\Delta W
\end{aligned}
\tag{16}
$$

Note that A and B are not dependent on c because dealers only learn of their hedging access after meeting with investors.

Lemma 2. ***(Prices and Bid-Ask Spread)*** *Bid and ask prices are linearly increasing in the dealer's ex-ante expected reservation value ΔW. Bid-ask spread is linearly decreasing in ΔW. The bid-ask spread is given by:*

$$
\begin{aligned}
A-B &= \theta(\Delta V_h-\Delta V_l) \\
&= \frac{r\lambda\theta(1-\theta)(\nu_l-\nu_n)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}\Delta W \\
&\quad+\frac{r\theta(\varepsilon_h-\varepsilon_l)+\lambda\theta(1-\theta)(\nu_n\varepsilon_h-\nu_l\varepsilon_l)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}
\end{aligned}
$$

Proof. See Appendix B.2 □

ΔW is dealer's ex-ante expected reservation value. It can be interpreted as dealers' willingness to provide liquidity. Notice that except for ΔW, no other term on the right-hand side depends on the hedging access, inter-dealer search intensity or dealer inventory holding cost. ΔW implicitly depends on these parameters. Dealers' ex-ante expected reservation value ΔW is increasing in the inter-dealer search intensity ρ, and the bid-ask spread is decreasing in ρ. That is, an increase in the inter-dealer market search frictions (i.e., decrease in search intensity) leads to lower ex-ante expected reservation value for dealers and higher bid-ask spread. Thus, higher search frictions in the inter-dealer market affect dealer liquidity provision and transaction costs for investors in an adverse manner. Furthermore, dealers' ex-ante expected reservation value ΔW is decreasing in the unhedged inventory cost c, and the bid-ask spread is thus increasing in c, when $\pi < 1$. In words, when the hedging technology is not always accessible, the higher the carrying cost the unhedged inventory is, the higher the bid-ask spread and the lower the liquidity in investor-dealer trading. Since c measures inventory frictions, the results says that when the hedging technology is not always accessible, inventory frictions lead to lower liquidity provision and higher transaction costs.[22]

It is not surprising that inter-dealer search frictions and dealer inventory frictions affect dealers' ex-ante expected reservation value when trading with investors, and hence the bid-ask spread. Search frictions in the inter-dealer market affect dealers' ability to offset inventory imbalances caused by transactions within investors. Inventory frictions are related to how burdensome inventory imbalances are to dealers. Rather, a key insight is that the OTC search frictions in the inter-dealer market interact with inventory frictions. The following proposition highlights the interdependence. In particular, search frictions and inventory frictions amplify each other in their effects on trading liquidity and transaction costs.

Proposition 3. ***(Inventory Frictions and Search)*** *The inter-dealer search intensity ρ interacts with the holding cost of unhedged inventory c in affecting dealers' ex-ante expected reservation value ΔW and the bid-ask spread. Mathematically,*

$$\frac{\partial^2 \Delta W}{\partial \rho \partial c} > 0$$

That is, as c increases, the magnitude of the effect of ρ on ΔW (and the bid-ask spread) increases. As ρ increases, the magnitude of the effect of c on ΔW (and the bid-ask spread) decreases.

Proof. See Appendix B.3 □

Proposition 3 says that when the inter-dealer search frictions are high (i.e., ρ is low), dealers' willingness to provide liquidity and transaction costs are more sensitive to inventory frictions, in the sense that the effect of inventory cost on liquidity provision and transaction costs is higher in magnitude. Conversely, when dealers face higher inventory frictions (i.e., c is high), dealers' willingness to provide liquidity and transaction costs are more sensitive to inter-dealer search frictions, in the sense that the effect of search frictions on liquidity provision and transaction costs is also higher in magnitude. These results are intuitive. Without a perfectly

[22]This is consistent with the empirical observation that more stringent regulations post-GFC targeting dealers' ability to commit capital are associated with lower dealer inventory levels and liquidity provisions, which are more pronounced during periods of stress (see Bao et al. (2018), Bessembinder et al. (2018), Dick-Nielsen and Rossi (2018)).

competitive and frictionless inter-dealer market to unload undesired positions, dealers are exposed to inventory imbalances. The more costly it is to hold these inventory positions, the larger the effect is of search frictions on dealers' willingness to provide liquidity and thus transaction costs. Conversely, when the search frictions are particularly high, the risk of inventory imbalances is more pronounced, amplifying the effect of inventory cost on dealer liquidity provision and transaction costs.

In the next proposition, I show how access to the inventory hedging technology affects dealer liquidity provision and the bid-ask spread.

Proposition 4. ***(Inventory Hedging and Liquidity)*** *Dealers' ex-ante expected reservation value ΔW is increasing in the accessibility of the hedging technology π. The bid-ask spread is decreasing in dealers' access to the hedging technology π. Mathematically,*

$$\frac{\partial \Delta W}{\partial \pi} > 0, \quad \frac{\partial (A - B)}{\partial \pi} < 0$$

Proof. See Appendix B.4 □

That is, when the hedging technology is accessible, dealers' willingness to provide liquidity is high and they make tighter market (lower bid-ask spread). Conversely, when their hedging options are lacking, dealers' liquidity provision is low and they make wider market (higher bid-ask spread). Intuitively, higher accessibility of the hedging technology decreases the expected inventory holding cost for dealers, thus increasing their ex-ante expected reservation value and hence their willingness to make market. Higher liquidity provision by dealers, in turn, improves trading liquidity and lowers transaction costs. Moreover, the effect of hedging availability on dealers' ex-ante expected reservation value and thus the bid-ask spread is nonlinear. Dealers' ex-ante expected reservation value ΔW is strictly concave in π, the accessibility of the hedging technology. Thus, the bid-ask spread is strictly convex in π. This means that dealer liquidity provision and transaction costs are more sensitive to changes in hedging accessibility π, when π is low.

Importantly, frictional inter-dealer market is critical for inventory hedging to have an impact on bid-ask spread. Consider the case with frictionless and competitive inter-dealer market. Given that the total supply of the tree s is small, any dealer who buys from investor and is subsequently revealed to not have access to the hedging technology can immediately unload its position to a nil dealer who can hedge. As a result, it is as if dealers always have access to the hedging technology, and the dealer sector can be reduced to a representative one. In this case, we are back to Duffie et al. (2005) and dealers' ex-ante expected reservation value and the bid-ask spread do not depend on the accessibility of the hedging technology. In the following proposition, I show that the effects of inventory hedging on dealer liquidity provision and transaction costs depend on the level of inventory frictions.

Proposition 5. ***(Hedging and Inventory Frictions)*** *The magnitude of the effect of hedging access π on dealers' ex-ante expected reservation value ΔW and the bid-ask spread is increasing in the unhedged inventory holding cost c. Mathematically,*

$$\frac{\partial^2 \Delta W}{\partial \pi \partial c} > 0, \quad \frac{\partial^2 (A-B)}{\partial \pi \partial c} < 0$$

Proof. See Appendix B.5 □

Proposition 5 says that when dealers face higher inventory frictions (i.e., unhedged inventory holding cost c is high), the impact of hedging access π on dealer liquidity provision and transaction costs is also high in magnitude. In words, inventory hedging is particularly beneficial when dealers face high inventory frictions. Cross-sectionally, riskier assets tend to impose higher inventory costs on dealers. Over time, dealers face higher inventory frictions when they face more intermediary constraints, such as higher funding costs, regulatory restrictions, and financial distress. Intuitively, when inventory frictions are high, the marginal benefit of hedging is also high, and thus the accessibility of the hedging technology tends to have an outsized impact on liquidity and transaction costs.

On the other hand, Proposition 5 also implies that when dealers lack hedging options (i.e., π is low), the effect of inventory cost c on dealers' willingness to provide liquidity and transaction cost is high in magnitude. That is, when hedging options are scarce, dealer liquidity provision and transaction costs are more sensitive to dealer inventory frictions, or changes in inventory cost. Conversely, when the hedging technology is accessible, liquidity and transaction costs have lower sensitivity to inventory frictions. This formalizes the intuition in Section 3.1 that hedged trades should exhibit lower sensitivity to inventory shocks.

4.3 Implications for Portfolio Trading

The model provides a conceptual framework to consider the mechanism of portfolio trading and how it affects liquidity and transaction costs. The underlying assumption is that baskets of corporate bonds can be more easily and cheaply hedged using credit index derivatives. First, diversification at the portfolio level reduces idiosyncratic risks, and thus hedging a portfolio of bonds is much easier by tracking key risk factors through exposures to indices. Second, credit index derivatives have experienced explosive growth over the past few years, and a highly liquid market has developed for standardized and centrally cleared credit index derivatives.

Thus, a direct interpretation under the conceptual framework is that portfolio trades can be more easily hedged with standardized index derivatives that are centrally cleared and highly liquid. Thus, portfolio trading is associated with a higher π under the conceptual framework. Alternatively, it could also be that portfolio trading results in a more hedgeable inventory, such as one that mimics an index and can be hedged with readily available index derivatives. In this sense, portfolio trading itself is a dealer inventory management tool, through which dealers buy or sell bonds needed in order to reduce the cost or increase the effectiveness of hedging their inventories. Again, this corresponds to a higher π in the model, and is fully consistent with the conceptual framework. Both mechanisms may be at work. Although it would be difficult to disentangle these two possible mechanisms, the second mechanism potentially explains why portfolio transactions consist of both dealer-buy and dealer-sell trades, even though dealers do not carry material short positions.

Since portfolio trading is associated with a higher π, or more accessible inventory hedging, in the model, the model yields the following testable predictions.

- **Prediction 1:** Transaction costs are lower for bonds traded as part of a portfolio transaction, relative to similar bonds traded through traditional voice trading. The weighted average transaction cost of a portfolio trade at the portfolio level is lower than that of a similar portfolio consisting of comparable voice-trade bonds.

- **Prediction 2:** Transaction cost advantage of bonds traded via portfolio trading is more pronounced for bonds associated with high inventory frictions (such as bonds with lower credit ratings, longer maturities and smaller issue sizes). Transaction cost advantage of portfolio-trade bonds is also larger when the dealer sector is more constrained.

5 Empirical Analysis

In this section, I test the model predictions laid out in Section 4.3. I show that, consistent with Prediction 1, portfolio trades have lower transaction costs at both the portfolio level and the individual bond level. Interesting, I document a cross venue effect, that is, a dealer's portfolio trading activity in a given bond seems to spillover into its voice trading, and reduces transaction costs of similar voice trades as well. I then attempt to disentangle various channels through which the cross venue effect takes place. Finally, consistent with Prediction 2, the transaction cost effect of portfolio trading is amplified for bonds with lower credit ratings, longer maturities and smaller issue sizes, and when the dealer sector is constrained.

5.1 Liquidity Benefit of Portfolio Trading

Prediction 1 of Section 4.3 says that portfolio trading reduces transaction costs, relative to similar voice-trade bonds. This is because portfolio trading benefits from more effective inventory hedging, which in turn improves liquidity and transaction costs. Table 8 illustrates the cost advantage of portfolio trading over voice trading for each of the four quarters in 2018 as well as the full estimation sample.

Table 8: Cost Advantage of Portfolio Trading

	Transaction Cost (bps)			Cost		
	Portfolio Trade	Matched Voice	All Voice	Advantage	t-statistic	*N*
By Line Item:						
Q1 2018	11.36	17.15	13.21	5.80	6.9	10,789
Q2 2018	7.93	14.27	13.90	6.34	12.0	19,801
Q3 2018	6.44	13.15	11.08	6.69	15.1	22,034
Q4 2018	8.08	15.07	12.78	6.96	19.5	39,212
Full Year	8.04	14.68	12.76	6.62	27.7	91,836
By Portfolio:						
Q1 2018	12.98	14.28		1.30	0.4	119
Q2 2018	6.81	13.41		6.59	3.9	183
Q3 2018	6.76	11.65		4.89	4.2	199
Q4 2018	9.99	13.94		3.96	2.7	319
Full Year	8.93	13.32		4.39	5.0	820

Note: This table reports the cost advantage of portfolio trades over voice trades for each of the four quarters in 2018 as well as full year 2018. Line items within portfolio trades are matched to comparable voice trades on the same day, by industry, rating, maturity, issue size and trade size. If there are multiple matches, the simple average across the matched transactions is obtained. To compare transaction costs at the portfolio level, the transaction cost of matched voice trades is the weighted average transaction costs of the matched voice trades, where the weights are the same weights in the portfolio trade.

I compare portfolio trades to voice trades both by line item and by portfolio. To facilitate comparison, I restrict my attention to the sub-sample containing portfolio-trading dealers only. This is because portfolio-trading dealers tend to be large core dealers, and the transaction costs of bonds traded by these dealers tend to be lower than periphery dealers. I first match each line item within portfolio trades to comparable voice trades on the same day. Transactions are matched by issuer industry, issue credit rating, years to maturity, issue size, and trade size. If there are multiple matches, I take the simple average across the matched transactions.[23] At the line

[23] 99.85% of all reported line items within portfolio trades are matched to comparable voice trades using this method.) do not have transaction cost estimates and are deleted from the estimation.

item level, I report the mean transaction costs for reported portfolio trade line items, and their corresponding matched voice trades. I also report the mean transaction costs of all voice trades during the same period.[24] Cost advantage is calculated by subtracting the mean portfolio trade transaction cost from the mean matched voice trade transaction cost. The t-statistics and number of observations for the calculated cost advantages are shown in the last two columns. Since portfolio trades price and transact securities at the portfolio-level, it is important to compare the aggregate portfolio-level transaction costs between portfolio trades and voice trades. To compute the benchmark voice-trade transaction costs for each portfolio trade, I calculate the weighted average transaction cost using the cost estimates of individual matched voice trades and applying the portfolio weights from the portfolio transaction. The cost advantage is again the mean difference between the transaction costs of portfolio trades at the portfolio-level and the imputed costs for the matched portfolios of voice trades.

The portfolio-level cost advantage is slightly lower than the cost advantage computed at the level of line items. This suggests that line items of smaller trade sizes likely benefit more from portfolio trading. The cost advantages of portfolio trading are substantial in magnitude, both by line item and by portfolio; and with the exception of the first quarter of 2018 when corporate bond portfolio trading was still in its infancy and Dealer A held a dominant market share of 77% of total portfolio trading volume, all cost advantage estimates have large t-statistics. The average portfolio-level transaction cost for portfolio trades in the estimation sample is 4.39 bps, which corresponds to a 33% cost advantage over a comparable portfolio of matched voice trades.

5.2 Cross Venue Effect

Dealers' voice trading and portfolio trading are not completely segmented, and thus it is reasonable to expect that a dealer's participation in portfolio trading also affects its market-making and liquidity in other trading venues, such as voice trading.[25] There are a few reasons for this spillover. First, portfolio trading may improve dealers' liquidity provision and pricing in voice trading due to information spillover and increased ability to source bonds. Second, dealers' voice trading may benefit indirectly from the ability to hedge inventories as a result of their portfolio trading activity. Third, as suggested by Hendershott and Madhavan (2015), endogenous venue selection by investors may also affect transaction costs in both portfolio trading and voice trading.

The above channels, however, manifest themselves differently in the data. From Table 3, bond characteristics are substantially similar for portfolio-trade bonds and voice-trade bonds, suggesting that the cost advantage of

[24]The transaction costs for the matched voice trades are lower than the mean transaction costs for all voice trades. This is likely attributable to larger trade sizes portfolio-trade bonds (see Table 3).

[25]For example, O'Hara and Zhou (2021b) study the cross venue effect of corporate bond electronic trading on voice trading.

portfolio trading over voice trading is not a result of easier-to-trade or ex-ante more liquid bonds migrating to portfolio trading due to investors' endogenous venue choice. Another possibility is that investors engaging in portfolio trading are more sophisticated or have higher bargaining power. Since these investors select into portfolio trading from voice trading, the cost advantage of portfolio trading also increases. However, since these investors migrate from voice trading after the advent of portfolio trading, the transaction costs of voice trading should increase as voice-trading investors are now on average less sophisticated. If, on the other hand, portfolio trading is associated with a decrease in voice trading transaction costs, then this channel is unlikely. In contrast, the first two channels both predict a decrease in voice trading transaction costs, associated with portfolio trading. However, if the main channel is the indirect effect of inventory hedging, then a dealer who engages in portfolio trading should also see its voice trades' inventory sensitivity decrease and bond turnover time increase.

5.2.1 Evidence of Cross Venue Effect

To formally test the cross venue effect, I first estimate the following specification:

$$Cost^{v}_{ijst} = \beta Traded\,Port_{ijst} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \tag{17}$$

where i indexes issue, j indexes dealer, s indexes trade size, and t indexes day. The outcome variable $Cost^{v}_{ijst}$ is the average transaction cost of a voice trade in bond i of size s executed by dealer j on day t. The explanatory variable $Traded\,Port_{ijst}$ is an indicator variable equal to one if the bond has traded as part of a portfolio trade in the same trade size bucket by the same dealer on the same day. Again, α_i, α_j, α_s and α_t are issue, dealer, trade size, and day fixed effects respectively. The regression estimate of specification (17) is reported in column (1) of Table 9. Column (2) saturates the specification with issue-size-day fixed effects, restricting comparison within the same bond of similar trade size on the same day, while controlling for differences across dealers. Both estimates are statistically significant at the 1% level. The results suggest that portfolio trading is associated with roughly 3-6 bps decrease in transaction costs of similar voice trades for portfolio-trading dealers.

I then test whether a dealer who trades more through portfolio trading has lower transaction costs in similar voice trades. I estimate the following specification:

$$Cost^{v}_{ijst} = \beta Port\,Share_{ijst} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \tag{18}$$

Table 9: Cross Venue Effect of Portfolio Trading on Voice Trading Transaction Costs

	Transaction Cost (bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
$Traded\ Port_{ijst}$	-5.938***	-3.441***				
	(1.686)	(1.026)				
$Port\ Share_{ijst}$			-0.100***	-0.041**		
			(0.032)	(0.017)		
$Post\ Port_{it}$					-0.884***	-1.052***
					(0.248)	(0.233)
Issue FE	Yes	No	Yes	No	Yes	No
Size FE	Yes	No	Yes	No	Yes	No
Day FE	Yes	No	Yes	No	Yes	No
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	Yes	No	Yes	No	Yes
R^2	0.115	0.395	0.115	0.395	0.115	0.395
Observations	4,574,555	3,214,167	4,574,555	3,214,167	4,574,555	3,214,167

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Columns (1) and (2) regress voice-trading transaction costs, $Cost^{v}_{ijst}$ on an indicator variable equal to one if the bond has traded as part of a portfolio trade in the same size bucket by the same dealer on the same day, while controlling for issue, trade size, day and dealer fixed effects, and issue-size-day, dealer fixed effects respectively. Columns (3) and (4) estimate the same specifications by replacing $Port\ Trade_{ijst}$ with the share of dealer j's trading volume in bond i of trade size s on day t that is transacted through portfolio trading. Columns (5) and (6) estimate a staggered difference-in-differences specification where voice-trading transaction costs are regressed on the indicator variable equal to one if bond i has been traded as part of a portfolio trade on or before day t. All standard errors are clustered at the issue-day and the dealer-day levels. The estimation sample covers corporate bond transaction data from the Academic TRACE between January 1, 2018 and December 31, 2018.

The explanatory variable $Port\ Share_{ijst}$ is the share of dealer j's total trading volume in bond i of size s on day t that is transacted through portfolio trading. If the cross venue effect is at work, then bonds for which a higher share of trading volume goes through portfolio trading should see more of a spillover effect in their voice trading. Columns (3) and (4) in Table 9 present the regression estimates for specification (18) and the saturated version with the interacted issue-size-day fixed effects, respectively. Both estimates are negative and highly statistically significant. To interpret the economic magnitude, note that $Port\ Share_{ijst}$ varies widely from zero to 100 percent, with a standard deviation of 13.5% in the estimation sample. Thus, a one standard deviation increase in a dealer's portfolio trading share in a given bond is associated with 0.6-1.3 bps decrease in the voice-trading transaction cost of the bond of similar trade size by the dealer on that day, or roughly 4-10% of the mean voice-trading transaction cost for portfolio-trading dealers in the estimation sample.

Finally, I test whether the voice-trading transaction cost of a bond decreases after the bond is traded in a portfolio trade. I estimate the following staggered difference-in-differences specification:

$$Cost^{v}_{ijst} = \beta Post\ Port_{it} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \tag{19}$$

where $Post\ Port_{it}$ is an indicator variable equal to one if bond i has been traded as part of a portfolio trade on or before day t. The fixed effects control for fixed differences across bond issues, dealers, trade sizes and time. In this difference-in-differences specification, the treatment is a bond having been ever traded as part of portfolio trading. Given that portfolio-trade bonds and voice-trade bonds are substantially similar across different bond characteristics, it is reasonable to assume parallel trends between the treated and the control. Again, I also saturate the specification by replacing α_i, α_s and α_t with α_{ist}, the interacted issue-size-day fixed effects. The resulting coefficient estimates are reported in columns (5) and (6) in Table 9 respectively. The coefficient estimates are negative and statistically significant at the 1% level, suggesting that once a bond has been traded as part of portfolio trading, transaction costs of similar voice trades tend to benefit.

In short, the above results support the presence of a cross venue effect between a dealer's portfolio trading and voice trading. That is, a dealer's portfolio trading activity appears to spillover to its voice trading and lowers the transaction costs of similar voice trades. These results are thus consistent with the first two channels hypothesized at the beginning of this section, instead of the endogenous venue selection channel. In the following section, I further distinguish between the first two channels, namely the search and information spillover versus the indirect effect of inventory hedging.

5.2.2 Potential Channels for Cross Venue Effect

As discussed, the cross venue effect may be due to search and information spillover or due to indirect effect of inventory hedging. To distinguish between the two channels, I test whether voice trades in bonds that are traded more through portfolio trading are less sensitive to bond flow shocks or stay in dealer inventories longer. If the cross venue effect is mainly driven by the indirect effect of inventory hedging, then these voice trades should exhibit similarly reduced inventory sensitivity and slower bond turnover as portfolio trades. To test whether portfolio trading leads to lower inventory exposure of similar voice trades, I estimate the following specification:

$$\begin{aligned} Cost^{v}_{ijst} =& \beta_1 Port\,Share_{ijst} + \beta_2 Flow\,Shock_{ijt} + \beta_3 Port\,Share_{ijst} \times Flow\,Shock_{ijt} \\ &+ \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \end{aligned} \tag{20}$$

where the outcome variable, the explanatory variables and the fixed effects are defined as before. The coefficient of interest is β_3. A negative β_3 estimate suggests that a larger portfolio trading share is associated with reduced sensitivity of similar voice-trades in the bond to inventory shocks. Thus, a negative β_3 estimate supports the potential explanation that the cross venue effect is due to indirect effect of inventory hedging associated with a dealer's portfolio trading activity. The results are presented in columns (1)-(3) in Table 10. Column (1) estimates the baseline specification (20), column (2) saturates the specification with issue-size-day fixed effects, and column (3) instruments $Flow\,Shock_{ijt}$ with the flow shock in matched bonds by the same dealer on the same day. None of the coefficient estimates is statistically significant at the conventional levels.

I then regress bond turnover, $Turnover_{ijkst}$ on $Port\,Share_{ijst}$, by estimating:

$$Turnover^{d}_{ijkst} = \beta Port\,Share_{ijst} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \tag{21}$$

and also saturate the specification by replacing the issue, size and day fixed effects with the interacted issue-size-day fixed effects. The coefficient estimates are reported in columns (4) and (5) of Table 10. Both coefficient estimates are negative and statistically significant at the 1% level, suggesting that portfolio trading is associated with faster bond turnover in similar voice trades. Hence, the results from Table 10 suggest that the cross venue effect is unlikely driven by the indirect effect of inventory hedging from portfolio trading activity.

Next, I show that the cross venue effect between portfolio trading and voice trading is likely due to search and information spillover. Since portfolio trading provides dealers an alternative way to source bonds, and

Table 10: Portfolio Trading and Inventory Exposure of Similar Voice Trades

	Transaction Cost (bps)			Bond Turnover (days)	
	(1)	(2)	(3)	(4)	(5)
$Port\,Share_{ijst} \times Flow\,Shock_{ijt}$	0.016	-0.001	0.040		
	(0.010)	(0.004)	(0.024)		
$Port\,Share_{ijst}$				-0.054***	-0.055***
				(0.021)	(0.015)
Issue FE	Yes	No	Yes	Yes	No
Size FE	Yes	No	Yes	Yes	No
Side FE	No	No	No	Yes	No
Day FE	Yes	No	Yes	Yes	No
Dealer FE	Yes	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	Yes	No	No	No
Issue-Size-Side-Day FE	No	No	No	No	Yes
R^2	0.116	0.396	-0.002	0.205	0.515
Observations	4,574,508	3,214,103	4,574,508	4,154,870	2,273,952
Specification	FE	FE	2SLS	FE	FE

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Columns (1)-(3) test whether portfolio trading is associated with less inventory exposure of similar voice trades. Column (1) regresses the voice-trading transaction costs, $Cost^{v}_{ijst}$, on the interaction between the portfolio-trading share $Port\,Share_{ijst}$ and bond flow shock $Flow\,Shock_{ijt}$ in the same bond by the same dealer on the same day, while controlling for issue, size, day and dealer fixed effects. Column (2) estimates the same specification, while controlling for the interacted issue-size-day, and dealer fixed effects. Column (3) estimates a 2SLS specification using flow shocks in matched bonds by the same dealer on the same day to instrument $Flow\,Shock_{ijt}$. Columns (4)-(6) test whether portfolio trading is associated with differential bond turnover time in similar voice trades. Column (4) regresses bond turnover measure on $Port\,Share_{ijst}$, while controlling for issue, trade size, trade side, day and dealer fixed effects. Column (5) estimates the same specification, but instead controlling for the interacted issue-size-side-day and dealer fixed effects. Column (6) repeats the regression in column (5), but replacing the explanatory variable with $Port\,Traded_{ijst}$, an indicator variable equal to one if bond i of size s is traded as part of a portfolio trade by dealer j on day t. Standard errors are double-clustered at the issue-day and the dealer-day levels. The estimation sample covers corporate bond transaction data from the Academic TRACE between January 1, 2018 and December 31, 2018.

because portfolio-trading dealers benefit from information gained from their portfolio-trading activity, these dealers encounter less search frictions in trading these bonds through traditional voice trading, resulting in lower voice-trading transaction costs. Following O'Hara and Zhou (2021b), I test whether portfolio trading is associated with lower agency trading share and lower inter-dealer trading share in similar voice trades. The agency trading share is the percentage of voice trades (in volume) that are pre-arranged, which I define as a roundtrip within 15 minutes between two trades in the same bond of opposite directions by the same dealer of the same trade size. Inter-dealer trading share is the percentage of a dealer's trading volume that is transacted with another dealer. If the search and information spillover channel is at work, a bond traded in portfolio trading should have a lower agency trading share, as the dealer is more willing to engage in principal trading in face of lower search frictions (see Goldstein and Hotchkiss (2020)), as well as a lower inter-dealer trading share, as the dealer is less reliant on the inter-dealer market to move bonds. I thus estimate the following baseline specification:

$$Y_{ijt} = \beta Port\,Share_{ijst} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijst} \tag{22}$$

where Y_{ijt} is either agency trading share or inter-dealer trading share in bond i by dealer j on day t. The results are reported in Table 11. Columns (1) and (3) estimate the baseline specification (22) with $Agency\,Share_{ijt}$ and $Inter\,Dealer\,Share_{ijt}$ as outcome variables respectively. Columns (2) and (4) saturate the specification with the interacted issue-size-day fixed effects. The coefficient estimates in columns (2)-(4) are all negative and statistically significant at the 1% level. The coefficient estimate in column (1) is negative and statistically significant at the 10% level. The results suggest that portfolio trading is indeed associated with a decrease in both agency trading share and inter-dealer trading share in similar voice trades, supporting the search and information spillover channel.

5.3 Interaction with Inventory Frictions

So far, I have shown that portfolio trading is associated with reduced transaction costs, consistent with Prediction 1 of Section 4.3. I also document a cross venue effect between portfolio trading and voice trading. Prediction 2 of Section 4.3 predicts that the liquidity benefit of portfolio trading is more pronounced for bonds associated with higher ex-ante inventory frictions, and when the dealer sector is more constrained. To test this, I interact the portfolio trade dummy with various bond characteristics as well as macro factors, and regress the average issue-size-dealer-day level transaction costs on these interaction terms, while controlling for various fixed

Table 11: Agency Trading Share and Inter-Dealer Trading Share

	Agency Share (%)		Inter-Dealer Share (%)	
	(1)	(2)	(3)	(4)
$Port\ Share_{ijst}$	-0.019*	-0.035***	-0.054***	-0.066***
	(0.010)	(0.011)	(0.008)	(0.010)
Issue FE	Yes	No	Yes	No
Size FE	Yes	No	Yes	No
Day FE	Yes	No	Yes	No
Dealer FE	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	Yes	No	Yes
R^2	0.545	0.682	0.496	0.617
Observations	4,582,234	3,219,559	4,582,234	3,219,559

Standard errors in parentheses

$^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: Columns (1) and (2) test whether portfolio-trading share affects the percentage of dealer-investor trading that is pre-arranged, defined as roundtrip trades within 15 minutes from each other. Column (1) regresses the agency trading share on the portfolio-trading share, $Port\ Share_{ijst}$, while controlling for issue, size, day and dealer fixed effects. Column (2) estimates the same specification, but instead controls for the interacted issue-size-day and dealer fixed effects. Columns (3) and (4) test whether portfolio-trading share affects the percentage of inter-dealer trading. Column (3) regresses the inter-dealer trading share on the portfolio-trading share, while controlling for issue, size, day and dealer fixed effects. Column (4) replaces the issue, size and day fixed effects in column (3) with the interacted issue-size-day fixed effects instead. Standard errors are double-clustered at the issue-day and dealer-day levels. The estimation sample covers corporate bond transaction data from the Academic TRACE between January 1, 2018 and December 31, 2018.

effects and macro controls. The basic specification is as follows:

$$\begin{aligned} Cost_{ijkst} =& \beta_1 Port\ Trade_{ijkst} + \beta_2 Friction\ Char_{it} + \beta_3 Port\ Trade_{ijkst} \times Friction\ Char_{it} \\ &+ \gamma X_t + Fixed\ Effects + \varepsilon_{ijkst} \end{aligned} \tag{23}$$

where $Friction\ Char_{it}$ are various bond-level, trade-level or macro variables that are related to intermediating frictions, including credit rating, years to maturity, log issue size, and the TED rate. In general, bonds with lower credit ratings, longer maturities and smaller issue sizes are associated with higher inventory frictions. These bonds tend to be riskier and there are more stringent risk management rules with respect to these securities within dealer firms (such as limitation on the amount of HY exposure). The TED rate is the spread between 3-month LIBOR and Treasury bills, and proxies the dealer sector's funding cost. When the dealer sector's funding cost is high, holding inventory can be particularly costly. To avoid multicollinearity, I include industry, rating, trade size, dealer and week fixed effects ($Fixed\ Effects$) as well as macro controls X_t. The results are presented in Table 12. Transaction costs of portfolio-trade bonds tend to benefit more relative to those of similar voice trades, when the ex-ante inventory frictions are higher, such as for bonds with lower credit ratings, longer maturities and smaller issue sizes. Furthermore, transaction costs of portfolio-trade bonds seem to benefit more when the dealer sector is more constrained, as proxied by high TED rate.

Table 12: Interaction with Issue, Trade and Macro Characteristics

	Transaction Cost (bps)			
	(1)	(2)	(3)	(4)
$Port\,Trade_{ijst}$ × Rating	-1.102***			
	(0.198)			
$Port\,Trade_{ijst}$ × Yrs to Maturity		-0.298**		
		(0.117)		
$Port\,Trade_{ijst}$ × log Issue Size			2.683***	
			(1.037)	
$Port\,Trade_{ijst}$ × TED Rate				-12.72*
				(7.054)
RatingFE	Yes	Yes	Yes	Yes
IndustryFE	Yes	Yes	Yes	Yes
SizeFE	Yes	Yes	Yes	Yes
WeekFE	Yes	Yes	Yes	Yes
DealerFE	Yes	Yes	Yes	Yes
R^2	0.022	0.022	0.021	0.022
Observations	4,585,083	4,600,713	4,600,713	4,600,713

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table regresses the transaction cost measure, $Cost_{ijst}$, the average transaction cost in bond i of trade size s executed by dealer j on day t, on the interactions between the portfolio trade dummy and several bond characteristics as well as the TED rate. The explanatory variables in columns (1)-(3) are the interaction between the portfolio trade dummy and issue credit rating, years to maturity and log issue size, respectively. All regressions control for rating, industry, trade size, week and dealer fixed effects. Standard errors are double-clustered at the issue-day and dealer-day levels. The estimation sample covers corporate bond transaction data from the Academic TRACE between January 1, 2018 and December 31, 2018.

6 Conclusion and Discussion

Corporate bond portfolio trading has experienced rapid growth since the end of 2017. The meteoric rise in portfolio trading is driven by more effective inventory hedging technology available to dealers. In particular, the proliferation of standardized and centrally cleared credit index derivatives has likely facilitated this transition. A stylized dynamic search model featuring frictional intermediation and inventory hedging illustrates the key intuition, and develops implications of inventory hedging for dealer liquidity provision and price formation.

Before this paper, the literature had little to say about dealer inventory management beyond setting prices and varying search intensities. The canonical OTC search models (e.g., Duffie et al. (2005), Lagos and Rocheteau (2009)) assume away dealer inventory under frictionless inter-dealer market, despite recent evidence that the inter-dealer markets are far from frictionless and dealer inventory frictions feature prominently in OTC markets.[26]

This paper suggests that inventory hedging plays an important role in dealer inventory management, which in turn has profound implications for market liquidity, transaction costs, and prices. Inventory hedging benefits liquidity and reduces transaction costs. Importantly, the liquidity benefit of inventory hedging is more pronounced when dealers face higher inventory frictions. Consistent with model predictions, portfolio trades are associated with lower transaction costs compared to similar voice trades, and the transaction cost advantage is more pronounced for bonds that impose higher inventory frictions and when the dealer sector is more constrained. Interestingly, a cross venue effect is present, that is, a dealer's portfolio trading benefits transaction costs in similar voice-trades. This cross venue benefit is related to search and information spillover from portfolio trading to voice trading.

6.1 Policy Implications

Recent studies by Krishnamurthy and Duffie (2016), Du et al. (2017), Fleckenstein and Longstaff (2020), Du and Schreger (2021) suggest that certain post-GFC regulatory reforms, particularly the supplementary leverage ratio (SLR), have imposed significant balance sheet space constraints on banks, causing banks to move away from intermediating safe assets. Haynes et al. (2018) further argue that the SLR decreases banks' ability in clearing derivatives as cash collaterals posted by customers increase dealers' total leverage exposure, even though they also make the derivatives clearing business usually low-risk. The result is a reduction in the availability of derivatives viable for inventory hedging purposes, leading to lower liquidity provision and intermediation in risky OTC markets as well as safe assets.

Furthermore, the paper highlights the interaction between OTC search frictions and inventory frictions, in the sense that these frictions amplify each other in their effects on liquidity. The results suggest that policies aimed at reducing dealer inventory constraints are particularly helpful for OTC assets experiencing high search frictions, for example during times of stress. Given that the Fed's unconventional monetary policies during the COVID-19 shock in March and April 2020 reduced dealers' funding costs and constraints on holding assets

[26]E.g., Green et al. (2007), Hollifield et al. (2017), Li and Schürhoff (2019), Friewald and Nagler (2019), Bessembinder et al. (2020), Goldstein and Hotchkiss (2020), He et al. (2021).

(e.g., O'Hara and Zhou (2021a)), the results from this paper are consistent with the observations by many that these policies were effective at improving liquidity and prices in a variety of OTC markets. The results may also help explain why the corporate credit facility (CCF) expansion on April 9, 2020, which included fallen angels and high-yield ETFs, had outsized impact on HY bonds relative to IG bonds.[27]

Lastly, studies (e.g., Bao et al. (2018), Bessembinder et al. (2018), Dick-Nielsen and Rossi (2018)) have found that the post-GFC regulatory reforms limit dealers' capital commitment, leading to a deterioration in market liquidity in OTC markets such as the corporate bond market, particularly during times of stress. On the other hand, this paper suggests that dealers endogenously respond to regulatory constraints through active inventory management and financial innovation.

[27] For example, Boyarchenko et al. (2020), D'Amico et al. (2020), Gilchrist et al. (2021), Haddad et al. (2021), Kargar et al. (2021), Nozawa and Qiu (2021) among others.

References

Afonso, Gara and Ricardo Lagos, "Trade Dynamics in the Market for Federal Funds," *Econometrica*, 2011, *83* (1), 263–313.

Amihud, Yakov, "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, 2002, *5* (1), 31–56.

_ **and Haim Mendelson**, "Dealership Market: Market-Making with Inventory," *Journal of Financial Economics*, 1980, *8* (1), 31–53.

Bao, Jack, Maureen O'Hara, and Xing (Alex) Zhou, "The Volcker Rule and Corporate Bond Market Making in Times of Stress," *Journal of Financial Economics*, 2018, *130* (1), 95–113.

Bessembinder, Hendrik, Chester Spatt, and Kumar Venkataraman, "A Survey of the Microstructure of Fixed-Income Markets," *Journal of Financial and Quantitative Analysis*, 2020, *55* (1), 1–45.

_ **, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman**, "Capital Commitment and Illiquidity in Corporate Bonds," *Journal of Finance*, 2018, *73* (4), 1615–1661.

_ **, William Maxwell, and Kumar Venkataraman**, "Market transparency, liquidity externalities, and institutional trading costs in corporate bonds," *Journal of Financial Economics*, 2006, *82* (2), 251–288.

Boyarchenko, Nina, Anna Kovner, and Or Shachar, "It's What You Say and What You Buy: A Holistic Evaluation of the Corporate Bond Facilities," Working Paper 2020.

Brunnermeier, Markus K. and Lasse Heje Pedersen, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 2008, *22* (6), 2201–2238.

Buti, Sabrina, Barbara Rindi, and Ingrid M. Werner, "Diving into Dark Pools," Working Paper 2022.

Cespa, Giovanni and Thierry Foucault, "Illiquidity Contagion and Liquidity Crashes," *Review of Financial Studies*, 2014, *27* (6), 1615–1660.

Chen, Jiakai, Haoyang Liu, Asani Sarkar, and Zhaogang Song, "Dealers and the Dealer of Last Resort: Evidence from MBS Markets in the COVID-19 Crisis," Federal Reserve Bank of New York Staff Reports, no. 933 2021.

Chernenko, Sergey and Adi Sunderam, “The Real Consequences of Market Segmentation,” *Review of Financial Studies*, 2012, *25* (7), 2041–2069.

Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, “The Determinants of Credit Spread Changes,” *Journal of Finance*, 2007, *56* (6), 2177–2207.

Comerton-Forde, Carole and Talis J. Putniņš, “Dark Trading and Price Discovery,” *Journal of Financial Economics*, 2015, *118* (1), 70–92.

D’Amico, Stefania, Vamsidhar Kurakula, and Stephen Lee, “Impacts of the Fed Corporate Credit Facilities through the Lenses of ETFs and CDX,” Working Paper 2020.

Dannhauser, Caitlin D. and Saeid Hoseinzade, “The Unintended Consequences of Corporate Bond ETFs: Evidence from the Taper Tantrum,” *Review of Financial Studies*, 2021.

Dick-Nielsen, Jens and Marco Rossi, “The Cost of Immediacy for Corporate Bonds,” *Review of Financial Studies*, 2018, *32* (1), 1–41.

_ and Thomas Kjaer Poulsen, “How to clean Academic TRACE data,” Working Paper 2019.

Drechsler, Itamar, Alan Moreira, and Alexi Savov, “Liquidity and Volatility,” Working Paper 2021.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, “Deviations from Covered Interest Rate Parity,” *Journal of Finance*, 2017, *73* (3), 915–957.

_ and Jesse Schreger, “CIP Deviations, the Dollar, and Frictions in International Capital Markets,” Working Paper 2021.

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen, “Over-the-Counter Markets,” *Econometrica*, 2005, *73* (6), 1815–1847.

_ , _ , and _ , “Valuation in Over-the-Counter Markets,” *Review of Financial Studies*, 2007, *20* (6), 1865–1900.

Ellul, Andrew, Chotibhak Jotikasthira, and Christian T. Lundblad, “Regulatory Pressure and Fire Sales in the Corporate Bond Market,” *Journal of Financial Economics*, 2011, *101* (3), 596–620.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, “The Emergence of Market Structure,” *Review of Economic Studies*, 2021, *forthcoming.*

__ , __ , **Guido Menzio, and Ursula Wiriadinata**, "Intermediation as Rent Extraction," Working Paper 2019.

Fleckenstein, Matthias and Francis A. Longstaff, "Renting Balance Sheet Space: Intermediary Balance Sheet Rental Costs and the Valuation of Derivatives," *Review of Financial Studies*, 2020, *33* (11), 5051–5091.

Fleming, Michael J. and Joshua V. Rosenberg, "How Do Treasury Dealers Manage Their Positions?," Working Paper 2008.

Friewald, Nils and Florian Nagler, "Over-the-Counter Market Frictions and Yield Spread Changes," *Journal of Finance*, 2019, *74* (6), 3217–3257.

Gilchrist, Simon, Bin Wei, Vivian Z. Sue, and Egon Zakrajšek, "The Fed Takes on Corporate Credit Risk: An Analysis of the Efficacy of the SMCCF," Working Paper 2021.

Glosten, Lawrence R. and Paul R. Milgrom, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 1985, *14* (1), 71–100.

Goldstein, Michael A. and Edith S. Hotchkiss, "Providing liquidity in an illiquid market: Dealer behavior in US corporate bonds," *Journal of Financial Economics*, 2020, *135* (1), 16–40.

Green, Richard C., Burton Hollifield, and Norman Schürhoff, "Financial Intermediation and the Costs of Trading in an Opaque Market," *Review of Financial Studies*, 2007, *20* (2), 275–314.

Gromb, Denis and Dimitri Vayanos, "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs," *Journal of Financial Economics*, 2002, *66* (2), 361–407.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

Haddad, Valentin, Alan Moreira, and Tyler Muir, "When Selling Becomes Viral: Disruptions in Debt Markets in the COVID-19 Crisis and the Fed's Response," *Review of Financial Studies*, 2021, *34* (11), 5309–5351.

Hasbrouck, Joel and Gideon Saar, "Low-Latency Trading," *Journal of Financial Markets*, 2013, *16* (4), 646–679.

Haynes, Richard, Lihong McPhail, and Haoxiang Zhu, "Assessing the Impact of the Basel III Leverage Ratio on the Competitive Landscape of US Derivatives Markets: Evidence from Options," CFTC Policy Brief 2018.

He, Zhiguo and Arvind Krishnamurthy, "A Model of Capital and Crises," *Review of Economic Studies*, 2012, *79* (2), 735–777.

_ and _ , "Intermediary Asset Pricing," *American Economic Review*, 2013, *103* (2), 732–770.

_ , Bryan Kelly, and Asaf Manela, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126* (1), 1–35.

_ , Paymon Khorrami, and Zhaogang Song, "Commonality in Credit Spread Changes: Dealer Inventory and Intermediary Distress," *Review of Financial Studies*, 2021, *forthcoming.*

Hendershott, Terrence and Ananth Madhavan, "Click or Call? Auction versus Search in the Over-the-Counter Market," *Journal of Finance*, 2015, *70* (1), 419–447.

Hill, Joanne M. and Frank J. Jones, "Equity Trading, Program Trading, Portfolio Insurance, Computer Trading and All That," *Financial Analysts Journal*, 1988, *44* (4), 29–38.

Ho, Thomas and Hans R. Stoll, "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial Economics*, 1981, *9* (1), 47–73.

Hollifield, Burton, Artem Neklyudov, and Chester Spatt, "Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations," *Review of Financial Studies*, 2017, *30* (9), 3048–3085.

Hong, Gwangheon and Arthur Warga, "An Empirical Study of Bond Market Transactions," *Financial Analyst Journal*, 2000, *56* (2), 32–46.

Hotchkiss, Edith S. and Tavy Ronen, "The Informational Efficiency of the Corporate Bond Market: An Intraday Analysis," *Review of Financial Studies*, 2015, *15* (5), 1325–1354.

Hugonnier, Julien, Benjamin Lester, and Pierre-Olivier Weill, "Frictional Intermediation in Over-the-Counter Markets," *Review of Economic Studies*, 2020, *87* (3), 1432–1469.

_ , _ , and _ , "Heterogeneity in Decentralized Asset Markets," *Theoretical Economics*, 2021, *forthcoming.*

Iercosan, Diana and Alexander Jiron, "The Value of Trading Relationships and Networks in the CDS Market," Working Paper 2017.

Junge, Benjamin and Anders B. Trolle, "Liquidity Risk in Credit Default Swap Markets," Working Paper 2013.

Kargar, Mahyar, Benjamin Lester, David Lindsay, Shuo Liu, Pierre-Olivier Weill, and Diego Zúñiga, "Corporate Bond Liquidity during the COVID-19 Crisis," *Review of Financial Studies*, 2021, *34* (11), 5352–5401.

Kozora, Matthew, Bruce Mizrach, Matthew Peppe, Or Shachar, and Jonathan Sokobin, "Alternative Trading Systems in the Corporate Bond Market," *Federal Reserve Bank of New York Staff Reports*, 2020, *no. 938.*

Krishnamurthy, Arvind and Darrell Duffie, "Passthrough Efficiency in the Fed's New Monetary Policy Setting," Kansas City Federal Reserve Symposium on Designing Resilient Monetary Policy Frameworks for the Future 2016.

Lagos, Ricardo and Guillaume Rocheteau, "Liquidity in Asset Markets with Search Frictions," *Econometrica*, 2009, *77* (2), 403–426.

Li, Dan and Norman Schürhoff, "Dealer Networks," *Journal of Finance*, 2019, *74* (1), 91–144.

Maggio, Marco Di, Amir Kermani, and Zhaogang Song, "The Value of Trading Relationships in Turbulent Times," *Journal of Financial Economics*, 2017, *124* (2), 266–284.

Merton, Robert C., "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 1974, *29* (2), 449–470.

Nagel, Stefan, "Evaporating Liquidity," *Review of Financial Studies*, 2012, *25* (7), 2005–2039.

Naik, Narayan Y. and Pradeep K. Yadav, "Do Dealer Firms Manage Inventory on a Stock-by-Stock or a Portfolio Basis?," *Journal of Financial Economics*, 2003, *69* (2), 325–353.

_ **and** _ , "Risk Management with Derivatives by Dealers and Market Quality in Government Bond Markets," *Journal of Finance*, 2003, *58* (5), 1873–1904.

Nozawa, Yoshio and Yancheng Qiu, "Corporate Bond Market Reactions to Quantitative Easing During the COVID-19 Pandemic," *Journal of Banking & Finance*, 2021, *133*, 106153.

O'Hara, Maureen and Xing (Alex) Zhou, "Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis," *Journal of Financial Economics*, 2021, *142* (1), 46–68.

_ and _ , "The electronic revolution of corporate bond dealers," *Journal of Financial Economics*, 2021, *140* (2), 368–390.

Paddrik, Mark and Stathis Tompaidis, "Market-Making Costs and Liquidity: Evidence from CDS Markets," Working Paper 2019.

Pan, Kevin and Yao Zeng, "ETF Arbitrage under Liquidity Mismatch," Working Paper 2019.

Reilly, Chris, "The Hidden Cost of Corporate Bond ETFs," Working Paper 2021.

Riggs, Lynn, Esen Onur, David Reiffen, and Haoxiang Zhu, "Swap Trading after Dodd-Frank: Evidence from Index CDS," *Journal of Financial Economics*, 2020, *137* (3), 857–886.

Schultz, Paul, "Corporate Bond Trading Costs: A Peek Behind the Curtain," *Journal of Finance*, 2002, *56* (2), 677–698.

Shim, John J. and Karamfil Todorov, "ETFs, Illiquid Assets, and Fire Sales," Working Paper 2021.

Siriwardane, Emil N., "Limited Investment Capital and Credit Spreads," *Journal of Finance*, 2019, *74* (5), 2303–2347.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo, "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics*, 2002, *20* (4), 518–529.

Stoll, Hans R., "The Supply of Dealer Services in Securities Markets," *Journal of Finance*, 1978, *33* (4), 1133–1151.

_ , "Portfolio trading," *Journal of Portfolio Management*, 1988, *14* (4), 20–24.

Üslü, Semih, "Pricing and Liquidity in Decentralized Asset Markets," *Econometrica*, 2019, *87* (6), 2079–2140.

Weill, Pierre-Olivier, "The Search Theory of OTC Markets," Working Paper 2021.

Yang, Ming and Yao Zeng, "Coordination and Fragility in Liquidity Provision," Working Paper 2021.

Appendix

A Figures and Tables

A.1 Figures

Figure A.1: iBoxx TRS Average Daily Volumes

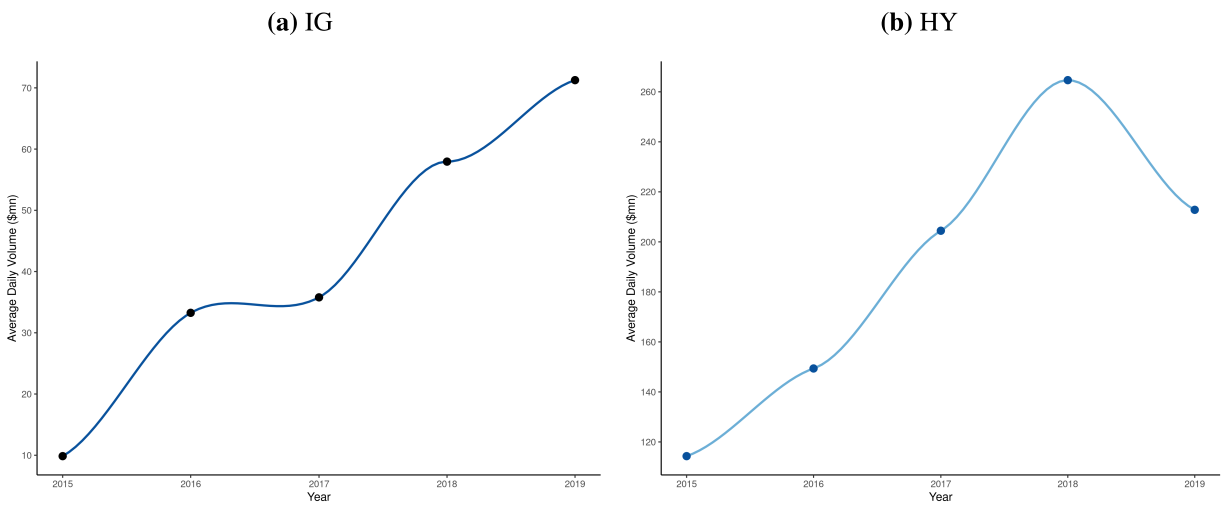

Note: This figure plots the average daily trading volumes in the iBoxx $ HY Index total return swap and the iBoxx $ IG Index total return swap respectively between 2015 and 2019. The trading volumes data are obtained from Markit website.

Figure A.2: Corporate Bond ETFs Listed Options Open Interest

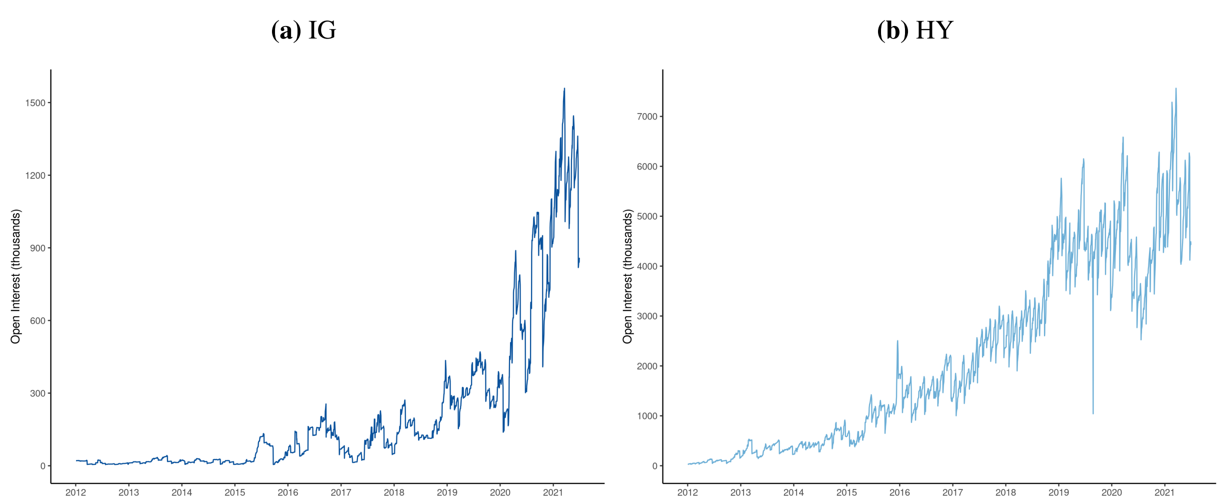

Note: This figure plots the daily open interest of listed options on LQD. Open interest data are obtained from Bloomberg.

Figure A.3: Growth of Corporate Bond ETFs

Note: The figure plots the daily time series of the net assets of the corporate bond ETFs against the total market value of the corporate bond market. An ETF's net assets are calculated as the product of its net asset value (NAV) per share and its total shares outstanding. ETF shares outstanding and NAV data are obtained from Bloomberg. The corporate bond market value is the combined market value of the Bloomberg Barclays U.S. Corporate Bond Index and the Bloomberg Barclays U.S. Corporate High Yield Bond Index. The index values are obtained from Bloomberg as well.

Figure A.4: Dealer-Customer TRACE Volume

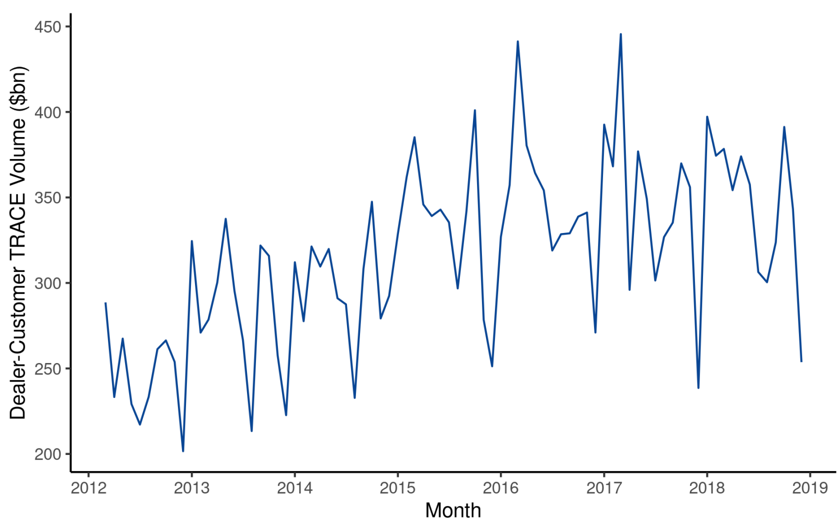

Note: The figure plots the monthly dealer-to-customer trading volume as reported by the Academic TRACE from March 1, 2018 to December 31, 2020.

Figure A.5: Portfolio Trading in the U.S. Corporate Bond Market Using Different Cutoffs

(a) Number of CUSIPs cutoff $N = 10$

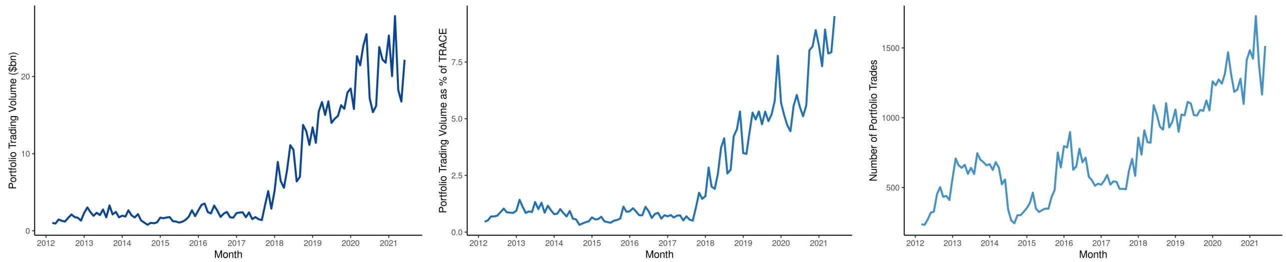

(b) Number of CUSIPs cutoff $N = 15$

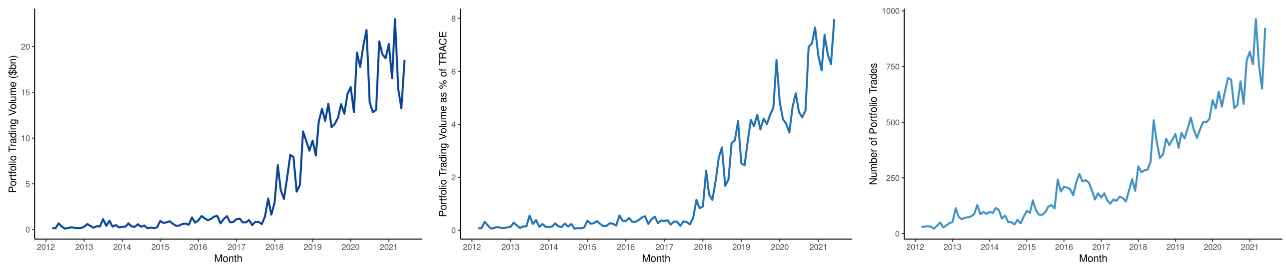

(c) Number of CUSIPs cutoff $N = 30$

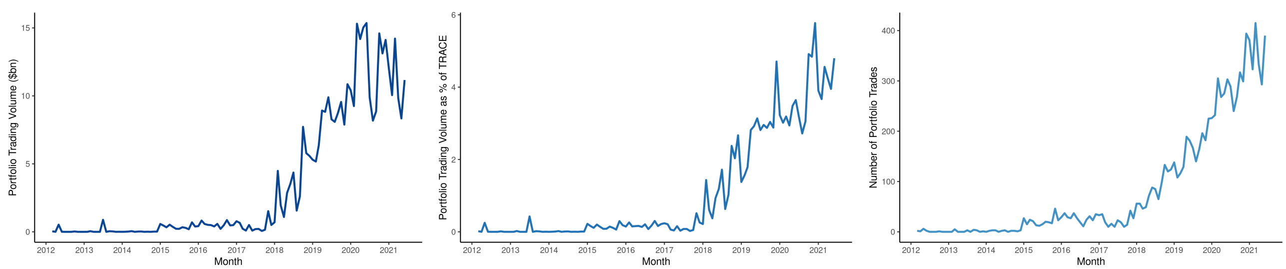

(d) Number of CUSIPs cutoff $N = 50$

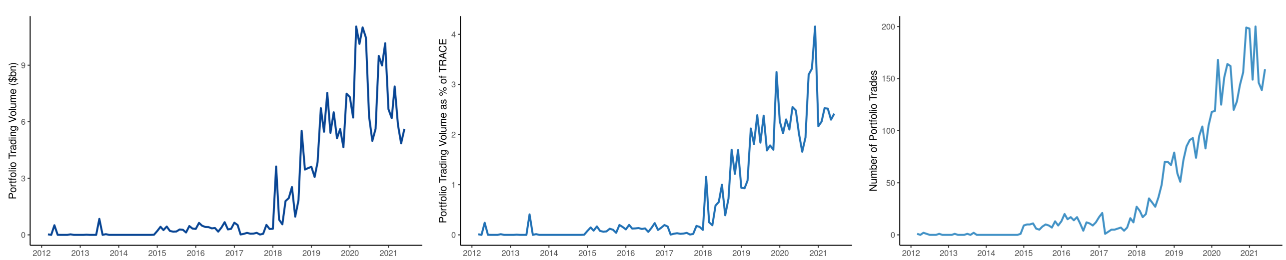

Note: This figure shows the rise of portfolio trading using different number-of-CUSIPs thresholds to identify portfolio trades. In each panel, the first figure from the left shows the monthly portfolio trading volume (in $bn) between March 2012 and June 2021. The second figure shows the monthly portfolio trading volume as a percentage of total dealer-to-customer trading volume according to TRACE during the same period. The third figure shows the number of portfolio trades on a monthly basis. Data are compiled from Enhanced TRACE available through the Wharton Research Data Services (WRDS).

A.2 Tables

Table A.1: Rating Scales

Moody's	S&P	Fitch	Value
Aaa	AAA	AAA	1
Aa1	AA+	AA+	2
Aa2	AA	AA	3
Aa3	AA-	AA-	4
A1	A+	A+	5
A2	A	A	6
A3	A-	A-	7
Baa1	BBB+	BBB+	8
Baa2	BBB	BBB	9
Baa3	BBB-	BBB-	10
Ba1	BB+	BB+	11
Ba2	BB	BB	12
Ba3	BB-	BB-	13
B1	B+	B+	14
B2	B	B	15
B3	B-	B-	16
Caa1	CCC+	CCC+	17
Caa2	CCC	CCC	18
Caa3	CCC-	CCC-	19
Ca	CC	CC	20
C	C	C	21
	D	DDD	22
		DD	23
		D	24

Note: This table reports the credit rating scales and associated numeric values assigned by the three major credit rating agencies, namely Moody's, S&P and Fitch. Credit ratings data for corporate bonds are obtained from the Mergent Fixed Income Securities Database (FISD). A numeric value of 10 (Baa3 by Moody's, BBB- by S&P and Fitch) or lower is considered investment-grade (IG) while a numeric value above 10 is considered non-investment grade or high-yield (HY). In the past, Moody's did not append numerical modifiers 1, 2, and 3 to each generic rating classification from Aa to Caa. To facilitate comparison, I map any generic rating to the mid-range ranking equivalent of the rating (for example, Aa is mapped to Aa2, and Baa is mapped to Baa2, etc.).

Table A.2: Differential Sensitivities of Transaction Costs to Bond Flow Shocks (Interactive Model)

	Transaction Cost (bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
$Port\,Trade_{ijkst} \times Flow\,Shock^{d}_{it}$	-1.421***			-1.015**		
	(0.526)			(0.514)		
$Port\,Trade_{ijkst} \times Flow\,Shock^{d}_{ijt}$		-4.210***	-2.940***		-3.451***	-2.856***
		(1.016)	(0.764)		(0.991)	(0.883)
Issue FE	Yes	Yes	No	Yes	Yes	No
Day FE	Yes	Yes	No	Yes	Yes	No
Size FE	Yes	Yes	No	Yes	Yes	No
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	No	Yes	No	No	Yes
R^2	0.115	0.115	0.396	0.063	0.063	0.493
Observations	4,664,572	4,664,572	3,292,736	1,948,452	1,948,452	877,847
Sample	Full	Full	Full	PTDealer	PTDealer	PTDealer

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table examines whether the transaction costs of portfolio-trade bonds and voice-trade bonds are differentially sensitive to bond flow shocks, by estimating the following main specification:

$$Cost_{ijkst} = \beta_1 Port\,Trade_{ijkst} + \beta_2 Flow\,Shock^{d}_{ijt} + \beta_3 Port\,Trade_{ijkst} \times Flow\,Shock^{d}_{ijt} + \alpha_i + \alpha_j + \alpha_s + \alpha_t + \varepsilon_{ijkst}$$

where i indexes issue, j indexes dealer, k indexes trade category (portfolio trade or voice trade), s indexes trade size, and t indexes time. $\alpha_i, \alpha_j, \alpha_s, \alpha_t$ are issue, dealer, trade size, and day fixed effects respectively. The coefficient of interest is β_3. $Flow\,Shock^{d}_{ijt}$ is an indicator variable equal to one if the issue-dealer-day level bond flow is above its sample median, and zero otherwise. $Port\,Trade_{ijkst}$ is the portfolio trade dummy. The outcome variable $Cost_{ijkst}$ is the average issue-dealer-size-day level transaction cost for type-k trade. The estimation sample covers U.S. corporate bond transaction data by dealers who engage in portfolio trading from January 1, 2018 to December 31, 2018. Column (1) estimates the above specification for the full-sample, but with the aggregate flow shock dummy $Flow\,Shock^{d}_{it}$ as the explanatory variable. Column (2) estimates the above specification for the full-sample, and column (3) saturates the specification by replacing the issue, trade size and day fixed effects with issue-size-day fixed effects. Columns (4)-(6) repeat the regressions in columns (1)-(3), but for the sub-sample containing only portfolio-trading dealers. All standard errors are double clustered at the issue-day and dealer-day levels.

Table A.3: Sensitivities of Transaction Costs to Bond Flow Shocks (Portfolio-Trading Dealers)

	Transaction Cost (bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
Voice-Trading:						
$Flow\,Shock_{it}$	0.312***	0.188***	0.105***			
	(0.037)	(0.025)	(0.020)			
$Flow\,Shock_{ijt}$				0.538***	0.095**	2.288**
				(0.032)	(0.037)	(0.960)
R^2	0.125	0.094	0.063	0.064	0.495	-0.004
Observations	955,207	1,384,288	1,858,423	1,858,423	811,474	1,858,423
Portfolio-Trading:						
$Flow\,Shock_{it}$	-0.052	-0.026	-0.014			
	(0.150)	(0.130)	(0.128)			
$Flow\,Shock_{ijt}$				-0.383	-0.399	-2.668
				(0.904)	(0.579)	(4.112)
R^2	0.230	0.227	0.230	0.230	0.617	-0.004
Observations	83,475	86,410	88,554	88,554	4,212	88,554
Issue FE	Yes	Yes	Yes	Yes	No	Yes
Day FE	Yes	Yes	Yes	Yes	No	Yes
Size FE	No	Yes	Yes	Yes	No	Yes
Dealer FE	No	No	Yes	Yes	Yes	Yes
Issue-Size-Day FE	No	No	No	No	Yes	No
Specification	FE	FE	FE	FE	FE	2SLS

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table reports the split-sample regression results examining the differential sensitivities in transaction costs of portfolio-trade bonds and voice-trade bonds to bond flow shocks. Bond flow shock is a dealer's net purchase of the bond on a given day. The estimation sample covers U.S. corporate bond transaction data by dealers who engage in portfolio trading from January 1, 2018 to December 31, 2018. The voice-trading sub-sample consists of investor-dealer trades not identified as part of a portfolio trade, while the portfolio-trading sub-sample includes all trades identified as part of a portfolio trade. The outcome variable in column (1) is $Cost_{it}$, the average transaction cost in bond i on day t. The outcome variable in column (2) is $Cost_{its}$, the average transaction cost in bond i of trade size s on day t. The outcome variable in column (3)-(6) is $Cost_{ijts}$, the average transaction cost of a trade with size s in bond i executed by dealer j on day t. The explanatory variable $Flow\,Shock_{ijt}$ is the net purchase of bond i by dealer j on day t (in \$mn). $Flow\,Shock_{it}$ is issue-dealer-day level flow shocks aggregated across all dealers (in \$mn). Column (5) further saturates the model with issue-size-day and dealer fixed effects. Column (6) instruments $Flow\,Shock_{ijt}$ with flow shock in comparable bonds matched by industry, rating, maturity and issue size categories, in a 2SLS specification. In columns (1) and (2), standard errors are clustered at the issue and the day levels. In columns (3)-(6), standard errors are double clustered at the issue-day and the dealer-day levels.

Table A.4: Impact of ETF-NAV Basis on Transaction Cost Differential

	PT Volume ($mn)		Transaction Cost (bps)			
	(1)	(2)	(3)	(4)	(5)	(6)
$ETF\,Arb\,Volume$	0.029					
	(0.025)					
$Net\,Create\,Volume$		0.030				
		(0.060)				
$Port\,Trade \times IG\,Basis$			2.076	-0.153		
			(13.84)	(11.32)		
$Port\,Trade \times HY\,Basis$					5.113	9.534
					(9.465)	(6.365)
Issue FE	No	No	Yes	No	Yes	No
Size FE	No	No	Yes	No	Yes	No
Day FE	No	No	Yes	No	Yes	No
Dealer FE	No	No	Yes	Yes	Yes	Yes
Issue-Size-Day	No	No	No	Yes	No	Yes
R^2	0.045	0.002	0.169	0.519	0.083	0.319
Observations	251	251	3,430,489	2,444,541	1,218,206	838,033
Sample	Full	Dlr-Buy	IG	IG	HY	HY

Note: Column (1) regresses daily portfolio trading volume on daily ETF arbitrage volume. Column (2) regresses daily portfolio trading volume involving dealer-buy trades on daily ETF net creation volume. ETF arbitrage volume is calculated as: $|\text{Shares Outstanding}_t - \text{Shares Outstanding}_{t-1}| \times NAV_t$, or the absolute value of the product of change in shares outstanding and NAV per share. ETF net creation volume is calculated according to the formula but only for ETFs whose shares outstanding have increased. Columns (3) and (4) regress transaction costs on the interaction between the portfolio trade dummy and LQD ETF-NAV basis for the IG sub-sample, while controlling for various fixed effects. Columns (5) and (6) regress transaction costs on the interaction between the portfolio trade dummy and HYG ETF-NAV basis for the HY sub-sample, while controlling for various fixed effects. ETF shares outstanding, share prices, and NAV per share data are both obtained from Bloomberg, while transaction cost measure is calculated according to Section 2.1.1. The estimation period is from January 1, 2018 to December 31, 2018. Standard errors for columns (3)-(6) are double clustered at the issue-day and dealer-day levels.

B Derivations and Proofs

B.1 Proof of Proposition 1

B.1.1 Convergence of μ_l and μ_h

Lemma. *Let $\mu_l(t) \equiv \mu_{lo}(t) + \mu_{ln}(t)$ denote the fraction of the investors who are low-type. Let $\mu_h(t) \equiv \mu_{ho}(t) + \mu_{hn}(t)$ denote the fraction of the investors who are high-type. Both $\mu_l(t)$ and $\mu_h(t)$ converge to a constant as $t \to \infty$, and their dynamics do not depend on μ_σ ($\sigma \in \{ho, hn, lo, ln\}$) and ν_ξ ($\xi \in \{l, n\}$).*

I first prove the above lemma. From (4), $\mu_h(t) = 1 - \mu_l(t)$. The dynamics for $\mu_{lo}(t)$ and $\mu_{ln}(t)$ are characterized in (7) and rewritten as follows:

$$\dot{\mu}_{lo}(t) = -\lambda \mu_{lo}(t)\nu_n(t) - \frac{1}{2}\mu_{lo}(t) + \frac{1}{2}\mu_{ho}(t) \tag{B.1}$$

$$\dot{\mu}_{ln}(t) = \lambda \mu_{lo}(t)\nu_n(t) - \frac{1}{2}\mu_{ln}(t) + \frac{1}{2}\mu_{hn}(t) \tag{B.2}$$

Adding up (B.1) and (B.2),

$$\begin{aligned}
\dot{\mu}_{lo}(t) + \dot{\mu}_{ln}(t) = \dot{\mu}_l(t) &= -\frac{1}{2}[\mu_{lo}(t) + \mu_{ln}(t)] + \frac{1}{2}[\mu_{ho}(t) + \mu_{hn}(t)] \\
&= -\frac{1}{2}\mu_l(t) + \frac{1}{2}\mu_h(t) \\
&= -\frac{1}{2}\mu_l(t) + \frac{1}{2}(1 - \mu_l(t)) \\
&= \frac{1}{2} - \mu_l(t)
\end{aligned}$$

The solution to the above ODE is

$$\mu_l(t) = \mu_l(0)e^{-t} + \frac{1}{2}\left(1 - e^{-t}\right)$$

which does not depend on other population measures. As $t \to \infty$,

$$\mu_l \to \frac{1}{2}$$

Since $\mu_h(t) = 1 - \mu_l(t)$,

$$\mu_h \to \frac{1}{2}$$

B.1.2 Existence and Uniqueness of Steady State

In a steady state equilibrium, the measure of each type of agents remains constant over time. Collect equations (4) to (6), suppressing the time argument t:

$$\mu_{ho} + \mu_{hn} + \mu_{lo} + \mu_{ln} = 1$$
$$\mu_{ho} + \mu_{lo} + \nu_l = s$$
$$\nu_l + \nu_n = 1$$

together with $\mu_{lo} + \mu_{ln} = \frac{1}{2}$ and $\mu_{ho} + \mu_{hn} = \frac{1}{2}$,

$$\mu_{ln} = \frac{1}{2} - \mu_{lo} \tag{B.3}$$
$$\mu_{ho} = \frac{1}{2} - \mu_{hn} \tag{B.4}$$
$$\nu_l = s - \mu_{ho} - \mu_{lo} = s - \frac{1}{2} + \mu_{hn} - \mu_{lo} \tag{B.5}$$

Note that in the dynamic system, $\dot{\nu}_l = 0$ and $\dot{\mu}_{hn} = 0$ imply that $\dot{\mu}_{ho} = \dot{\mu}_{lo} = \dot{\mu}_{ln} = \dot{\nu}_n = 0$. Set $\dot{\nu}_l$ to zero, and use $\nu_l + \nu_n = 1$, then

$$-\lambda \mu_{hn} \nu_l + \lambda \mu_{lo} \nu_n = 0 \iff \quad \nu_l = \frac{\mu_{lo}}{\mu_{lo} + \mu_{hn}} \tag{B.6}$$

From (B.5) and (B.6),

$$s - \frac{1}{2} + \mu_{hn} - \mu_{lo} = \frac{\mu_{lo}}{\mu_{lo} + \mu_{hn}}$$

Rearranging,

$$\mu_{hn}^2 + \left(s - \frac{1}{2}\right)\mu_{hn} - \left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right] = 0 \tag{B.7}$$

Solve for ν_{hn} in terms of ν_{lo} and restrict to the positive root:

$$\mu_{hn} = \frac{\frac{1}{2} - s + \sqrt{\left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right]}}{2} \tag{B.8}$$

Substitute (B.3), (B.5) and (B.8) into the expression for $\dot{\mu}_{hn}$:

$$\begin{aligned}
\dot{\mu}_{hn} &= -\lambda \mu_{hn}\left(s - \frac{1}{2} + \mu_{hn} - \mu_{lo}\right) - \frac{1}{2}\mu_{hn} + \frac{1}{2}\left(\frac{1}{2} - \mu_{lo}\right) \\
&= -\lambda\left[\mu_{hn}^2 + \left(s - \frac{1}{2}\right)\mu_{hn}\right] + \lambda \mu_{hn}\mu_{lo} - \frac{1}{2}\mu_{hn} + \frac{1}{2}\left(\frac{1}{2} - \mu_{lo}\right) \\
&= -\lambda\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right] + \lambda \mu_{hn}\mu_{lo} - \frac{1}{2}\mu_{hn} + \frac{1}{2}\left(\frac{1}{2} - \mu_{lo}\right) \\
&= -\lambda \mu_{lo}\left\{\mu_{lo} + \frac{3}{2} - s - \frac{1}{2}\left[\frac{1}{2} - s + \sqrt{\left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right]}\right]\right\} \\
&\quad - \frac{1}{4}\left[\frac{1}{2} - s + \sqrt{\left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right]}\right] + \frac{1}{2}\left(\frac{1}{2} - \mu_{lo}\right)
\end{aligned}$$

The third equality follows from (B.7). Notice that the last two terms are both decreasing in μ_{lo}. Next, I show that the first term is also decreasing in μ_{lo}. This is equivalent to showing that

$$h(\mu_{lo}) \equiv \mu_{lo}\left\{\mu_{lo} + \frac{3}{2} - s - \frac{1}{2}\left[\frac{1}{2} - s + \sqrt{\left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right]}\right]\right\}$$

is increasing in μ_{lo}. Since $s < \frac{1}{2}$,

$$\left(2\mu_{lo} + \frac{1}{2} - s\right)^2 \leq \left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right] \leq \left(2\mu_{lo} + \frac{3}{2} - s\right)^2 \tag{B.9}$$

Take derivative of $h(\mu_{lo})$ with respect to μ_{lo}:

$$\begin{aligned}
\frac{\partial h(\mu_{lo})}{\partial \mu_{lo}} &= \mu_{lo} + \frac{3}{2} - s - \frac{1}{2}\left[\frac{1}{2} - s + \sqrt{\left(\frac{1}{2} - s\right)^2 + 4\left[\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right]}\right] \\
&\quad + \mu_{lo}\left\{1 - \left[\left(\frac{1}{2} - s\right)^2 + 4\left(\mu_{lo}^2 + \left(\frac{3}{2} - s\right)\mu_{lo}\right)\right]^{-\frac{1}{2}}\left(2\mu_{lo} + \frac{3}{2} - s\right)\right\} \\
&\geq \mu_{lo} + \frac{3}{2} - s - \frac{1}{2}\left[2\mu_{lo} + 2 - 2s\right] + \mu_{lo}\left[1 - \frac{2\mu_{lo} + \frac{3}{2} - s}{2\mu_{lo} + \frac{1}{2} - s}\right] \\
&= \frac{1}{2} - \frac{\mu_{lo}}{2\mu_{lo} + \frac{1}{2} - s} > 0
\end{aligned}$$

Thus, I have shown that $\dot{\mu}_{hn}$ is decreasing in μ_{lo}. Let

$$
\begin{aligned}
Q(x) = -\lambda x\Big\{x + \frac{3}{2} - s - \frac{1}{2}\Big[\frac{1}{2} - s + \sqrt{\Big(\frac{1}{2} - s\Big)^2 + 4\Big[x^2 + \Big(\frac{3}{2} - s\Big)x\Big]}\Big]\Big\} \\
- \frac{1}{4}\Big[\frac{1}{2} - s + \sqrt{\Big(\frac{1}{2} - s\Big)^2 + 4\Big[x^2 + \Big(\frac{3}{2} - s\Big)x\Big]}\Big] + \frac{1}{2}\Big(\frac{1}{2} - x\Big)
\end{aligned}
$$

Thus,

$$
\begin{aligned}
Q(s) &= -\lambda s\Big\{\frac{3}{2} - \frac{1}{2}\Big[\frac{1}{2} - s + \sqrt{\Big(\frac{1}{2} - s\Big)^2 + 6s}\Big]\Big\} - \frac{1}{4}\Big[\frac{1}{2} - s + \sqrt{\Big(\frac{1}{2} - s\Big)^2 + 6s}\Big] + \frac{1}{2}\Big(\frac{1}{2} - s\Big) \\
&= -\frac{3}{2}\lambda s + \frac{1}{2}\Big(\lambda s - \frac{1}{2}\Big)\Big[\frac{1}{2} - s + \sqrt{\Big(\frac{1}{2} - s\Big)^2 + 6s}\Big] + \frac{1}{4} - \frac{1}{2}s
\end{aligned}
$$

If $\lambda s \geq \frac{1}{2}$,

$$
\begin{aligned}
Q(s) &\leq -\frac{3}{2}\lambda s + \frac{1}{2}\Big(\lambda s - \frac{1}{2}\Big)\Big[\frac{1}{2} - s + \frac{3}{2} + s\Big] + \frac{1}{4} - \frac{1}{2}s \\
&= -\frac{1}{2}\lambda s - \frac{1}{4} - \frac{1}{2}s < 0
\end{aligned}
$$

If $\lambda s < \frac{1}{2}$,

$$
\begin{aligned}
Q(s) &\leq -\frac{3}{2}\lambda s + \frac{1}{2}\Big(\lambda s - \frac{1}{2}\Big)\Big[\frac{1}{2} - s + \frac{1}{2} + s\Big] + \frac{1}{4} - \frac{1}{2}s \\
&= -\lambda s - \frac{1}{2}s < 0
\end{aligned}
$$

Thus, $Q(s) < 0$. Similarly, it is easy to show that $Q(0) > 0$. Therefore, $Q(x) = 0$ has a unique solution lying between 0 and s, that is, $\bar{\mu}_{lo} \in (0, s)$. Since $\bar{\mu}_{lo} + \bar{\mu}_{ln} = \frac{1}{2}$, it is immediate that $\bar{\mu}_{ln} \in \Big(0, \frac{1}{2}\Big)$.

From (B.8), it is obvious that μ_{hn} is increasing in μ_{lo}. Thus $\dot{\mu}_{hn}$ is decreasing in μ_{hn}. Express $\dot{\mu}_{hn}$ in terms of μ_{hn}:

$$
\dot{\mu}_{hn} = -\lambda\mu_{hn}\Big(s - \frac{1}{2} + h_{hn}\Big) + \Big(\lambda\mu_{hn} - \frac{1}{2}\Big)\frac{-\Big(\frac{3}{2} - s\Big) + \sqrt{\Big(\frac{3}{2} - s\Big)^2 + 4\mu_{hn}^2 + 4\Big(s - \frac{1}{2}\Big)\mu_{hn}}}{2} - \frac{1}{2}\mu_{hn} + \frac{1}{4}
$$

Let

$$\tilde{Q}(x) \equiv -\lambda x\left(s-\frac{1}{2}+x\right)+\left(\lambda x-\frac{1}{2}\right)\frac{s-\frac{3}{2}+\sqrt{\left(\frac{3}{2}-s\right)^2+4x^2+4\left(s-\frac{1}{2}\right)x}}{2}-\frac{1}{2}x+\frac{1}{4}$$

It is immediate that

$$\tilde{Q}(0)=\frac{1}{4}>0$$

and

$$\tilde{Q}\left(\frac{1}{2}\right)=-\frac{1}{2}\lambda s+\frac{1}{4}(\lambda-1)\left[s-\frac{3}{2}+\sqrt{\left(\frac{3}{2}-s\right)^2+2s}\right]$$

If $\lambda \geq 1$,

$$\tilde{Q}\left(\frac{1}{2}\right)\leq -\frac{1}{2}\lambda s+\frac{1}{2}(\lambda-1)s=-\frac{1}{2}s<0$$

If $\lambda < 1$,

$$\tilde{Q}\left(\frac{1}{2}\right)\leq -\frac{1}{2}\lambda s<0$$

Further,

$$\tilde{Q}\left(\frac{1}{2}-s\right)=-\frac{1}{2}\left(\frac{1}{2}-s\right)+\frac{1}{4}=\frac{1}{2}s>0$$

Hence, $\bar{\mu}_{hn} \in \left(\frac{1}{2}-s, \frac{1}{2}\right)$. Thus, $\bar{\mu}_{ho} \in (0,s)$. It is immediate from (B.6) that $\bar{v}_l > 0$ and with (B.5), we can also verify that $\bar{v}_l < s$.

B.1.3 Stability of Steady State

Denote the solution above as $(\bar{\mu}_{ho}, \bar{\mu}_{hn}, \bar{\mu}_{lo}, \bar{\mu}_{ln}, \bar{\nu}_l, \bar{\nu}_n)$. To show stability, notice that since

$$\mu_{ho} + \mu_{hn} + \mu_{lo} + \mu_{ln} = 1$$
$$\mu_{ho} + \mu_{lo} + \nu_l = s$$
$$\mu_{lo} + \mu_{ln} = \mu_l$$

we can express μ_{ho}, μ_{lo} and μ_{ln} in terms of μ_{hn}, ν_l and μ_l:

$$\mu_{ln} = 1 - s + \nu_l - \mu_{hn}$$
$$\mu_{lo} = \mu_l + s - 1 - \nu_l + \mu_{hn}$$
$$\mu_{ho} = 1 - \mu_l - \mu_{hn}$$

Thus, rewrite $\dot{\mu}_{hn}$ and $\dot{\nu}_l$ as:

$$\dot{\mu}_{hn} = -\lambda \mu_{hn}\nu_l - \mu_{hn} + \frac{1}{2}\nu_l + \frac{1}{2}(1-s)$$
$$\dot{\nu}_l = -2\lambda \mu_{hn}\nu_l + \lambda \nu_l^2 + \lambda \mu_{hn} - \lambda(\mu_l + s)\nu_l + \lambda(\mu_l + s - 1)$$

From Lemma 2, it is immediate that with large enough t_1, $\mu_l(t)$ changes by at most an arbitrarily small amount $\zeta > 0$ for all $t \geq t_1$, and it does not depend on μ_{hn} or ν_l. We also know that $\dot{\mu}_{hn} = 0$ and $\dot{\nu}_l = 0$ imply that $\dot{\mu}_{ho} = \dot{\mu}_{lo} = \dot{\mu}_{ln} = \dot{\nu}_n = 0$. I follow Lyapunov's indirect method. Let $\mathbf{x} \equiv (\mu_{hn}, \nu_l)$, and denote the system as $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ where $\mathbf{f} : \mathbb{R}^2 \to \mathbb{R}^2$. The Jacobian matrix associated with the linearization around the steady state $(\bar{\mu}_{hn}, \bar{\nu}_l)$ is:

$$\mathbf{A} = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}\Big|_{\mathbf{x}=(\bar{\mu}_{hn},\bar{\nu}_l)} = \begin{bmatrix} -(\lambda \bar{\nu}_l + 1) & -\lambda \bar{\mu}_{hn} + \frac{1}{2} \\ -2\lambda \bar{\nu}_l + \lambda & -2\lambda \bar{\mu}_{hn} + 2\lambda \bar{\nu}_l - \lambda(\bar{\mu}_l + s) \end{bmatrix}$$

If the system $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ is asymptotically stable, then the equilibrium point above is locally asymptotically stable. The eigenvalues of $\mathbf{A}$ are

$$\lambda_{\pm} = \frac{tr(\mathbf{A}) \pm \sqrt{tr(\mathbf{A})^2 - 4det(\mathbf{A})}}{2}$$

where

$$
\begin{aligned}
tr(\mathbf{A}) &= -\lambda \bar{\nu}_l - 1 - \lambda \left[2\bar{\mu}_{hn} - 2\bar{\nu}_l + \bar{\mu}_l + s\right] \\
&= -\lambda \bar{\nu}_l - 1 - \lambda \left(\frac{3}{2} - s + 2\bar{\mu}_{lo}\right) < 0
\end{aligned}
$$

and

$$
\begin{aligned}
det(\mathbf{A}) &= \lambda(\lambda \bar{\nu}_l + 1)(2\bar{\mu}_{hn} - 2\bar{\nu}_l + \bar{\mu}_l + s) - \lambda \left(\frac{1}{2} - \lambda \bar{\mu}_{hn}\right)(1 - 2\bar{\nu}_l) \\
&= \lambda \left[1 - s + 2\bar{\mu}_{lo} + \lambda \bar{\nu}_l \left(\frac{1}{2} + s - 2\bar{\nu}_l\right) + \bar{\nu}_l + \lambda \bar{\mu}_{hn}\right] \\
&> 0
\end{aligned}
$$

Therefore, all eigenvalues are in the open left half plane (OLHP). Hence, the equilibrium point is locally asymptotically stable.

B.2 Proof of Lemma 2

From (16), rearranging the first two equations

$$
\begin{aligned}
\left(r + \frac{1}{2} + \lambda \nu_n\right)\Delta V_l &= \frac{1}{2}\Delta V_h + \lambda \nu_n B + \varepsilon_l \\
\left(r + \frac{1}{2} + \lambda \nu_l\right)\Delta V_h &= \frac{1}{2}\Delta V_l + \lambda \nu_l A + \varepsilon_h
\end{aligned}
$$

Substitute the expressions of A and B into the above,

$$
\begin{aligned}
\left(r + \frac{1}{2} + \lambda \nu_n\right)\Delta V_l &= \frac{1}{2}\Delta V_h + \lambda \nu_n \theta \Delta V_l + \lambda \nu_n (1-\theta)\Delta W + \varepsilon_l \\
\left(r + \frac{1}{2} + \lambda \nu_l\right)\Delta V_h &= \frac{1}{2}\Delta V_l + \lambda \nu_l \theta \Delta V_h + \lambda \nu_l (1-\theta)\Delta W + \varepsilon_h
\end{aligned}
$$

Solve for ΔV_l and ΔV_h in terms of ΔW, we get

$$
\Delta V_l = a_1 \Delta W + b_1
$$

$$
\Delta V_h = a_2 \Delta W + b_2
$$

where

$$a_1 = \frac{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))\lambda\nu_n(1-\theta)+\frac{1}{2}\lambda\nu_l(1-\theta)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}$$
$$b_1 = \frac{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))\varepsilon_l+\frac{1}{2}\varepsilon_h}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}$$
$$a_2 = \frac{(r+\frac{1}{2}+\lambda\nu_n(1-\theta))\lambda\nu_l(1-\theta)+\frac{1}{2}\lambda\nu_n(1-\theta)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}$$
$$b_2 = \frac{(r+\frac{1}{2}+\lambda\nu_n(1-\theta))\varepsilon_h+\frac{1}{2}\varepsilon_l}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}$$

Hence,

$$A = \theta\Delta V_h + (1-\theta)\Delta W = (\theta a_2 + 1 - \theta)\Delta W + \theta b_2 \tag{B.10}$$

$$B = \theta\Delta V_l + (1-\theta)\Delta W = (\theta a_1 + 1 - \theta)\Delta W + \theta b_1 \tag{B.11}$$

Since $a_1, a_2 > 0$, both A and B are increasing in ΔW. Furthermore, notice that a_1, a_2, b_1, b_2 do not depend on the accessibility of the hedging technology π, inter-dealer market friction ρ or inventory cost c.

Bid-ask spread is thus

$$\begin{aligned} A - B &= \theta(\Delta V_h - \Delta V_l) \\ &= \frac{r\lambda\theta(1-\theta)(\nu_l-\nu_n)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}}\Delta W \\ &\quad + \frac{r\theta(\varepsilon_h-\varepsilon_l)+\lambda\theta(1-\theta)(\nu_n\varepsilon_h-\nu_l\varepsilon_l)}{(r+\frac{1}{2}+\lambda\nu_l(1-\theta))(r+\frac{1}{2}+\lambda\nu_n(1-\theta))-\frac{1}{4}} \end{aligned}$$

Since $\nu_l < s < \nu_n$, the bid-ask spread is linearly decreasing in ΔW.

B.3 Proof of Proposition 3 and Related Derivations

B.3.1 Derivation of ΔW

Rearrange (14):

$$
\begin{aligned}
[r+\lambda(\mu_{hn}+\mu_{lo})]W_l^{c_L} &= \lambda\mu_{hn}(A-\Delta W)+\lambda(\mu_{hn}+\mu_{lo})W_l \\
[r+\lambda(\mu_{hn}+\mu_{lo})]W_l^{c_H} &= \lambda\mu_{hn}(A-\Delta W)+\lambda(\mu_{hn}+\mu_{lo})W_l+\rho\nu_n\pi(\Delta W^{c_L}-\Delta W^{c_H})-c \\
[r+\lambda(\mu_{hn}+\mu_{lo})]W_n^{c_L} &= \lambda\mu_{lo}(\Delta W-B)+\lambda(\mu_{hn}+\mu_{lo})W_n+\rho\nu_l(1-\pi)(\Delta W^{c_L}-\Delta W^{c_H}) \\
[r+\lambda(\mu_{hn}+\mu_{lo})]W_n^{c_H} &= \lambda\mu_{lo}(\Delta W-B)+\lambda(\mu_{hn}+\mu_{lo})W_n
\end{aligned}
\tag{B.12}
$$

Thus,

$$
\begin{aligned}
W_l^{c_L}-W_l^{c_H} &= \frac{1}{r+\lambda(\mu_{hn}+\mu_{lo})}\left[c-\rho\nu_n\pi(\Delta W^{c_L}-\Delta W^{c_H})\right] \\
W_n^{c_L}-W_n^{c_H} &= \frac{1}{r+\lambda(\mu_{hn}+\mu_{lo})}\left[\rho\nu_l(1-\pi)(\Delta W^{c_L}-\Delta W^{c_H})\right]
\end{aligned}
$$

Subtract the second equation from the first,

$$
\Delta W^{c_L}-\Delta W^{c_H} = \frac{1}{r+\lambda(\mu_{hn}+\mu_{lo})}\left[c-\rho\nu_n\pi(\Delta W^{c_L}-\Delta W^{c_H})-\rho\nu_l(1-\pi)(\Delta W^{c_L}-\Delta W^{c_H})\right]
$$

Solve for $\Delta W^{c_L}-\Delta W^{c_H}$, we get

$$
\Delta W^{c_L}-\Delta W^{c_H} = \frac{c}{r+\lambda(\mu_{hn}+\mu_{lo})+\rho\nu_n\pi+\rho\nu_l(1-\pi)} \tag{B.13}
$$

To simplify notation, denote $\phi\equiv\Delta W^{c_L}-\Delta W^{c_H}$. From (B.10) and (B.11)

$$
\begin{aligned}
A-\Delta W &= \theta\left[(a_2-1)\Delta W+b_2\right] \\
\Delta W-B &= \theta\left[(1-a_1)\Delta W-b_1\right]
\end{aligned}
$$

From (B.12),

$$\begin{aligned}
\Delta W^{c_H} &= W_l^{c_H} - W_n^{c_H} \\
&= \frac{1}{r+\lambda(\mu_{hn}+\mu_{lo})}\Big[\lambda\mu_{hn}(A-\Delta W) - \lambda\mu_{lo}(\Delta W - B) + \lambda(\mu_{hn}+\mu_{lo})\Delta W + \rho\nu_n\pi\phi - c\Big] \\
\Delta W^{c_L} &= W_l^{c_L} - W_n^{c_L} \\
&= \frac{1}{r+\lambda(\mu_{hn}+\mu_{lo})}\Big[\lambda\mu_{hn}(A-\Delta W) - \lambda\mu_{lo}(\Delta W - B) + \lambda(\mu_{hn}+\mu_{lo})\Delta W - \rho\nu_l(1-\pi)\phi\Big]
\end{aligned}$$

Since $\Delta W = \pi\Delta W^{c_L} + (1-\pi)\Delta W^{c_H}$,

$$\begin{aligned}
[r+\lambda(\mu_{hn}+\mu_{lo})]\Delta W &= \lambda\mu_{hn}(A-\Delta W) - \lambda\mu_{lo}(\Delta W - B) + \lambda(\mu_{hn}+\mu_{lo})\Delta W \\
&\quad + \rho\pi(1-\pi)\phi\nu_n - \rho\pi(1-\pi)\phi\nu_l - (1-\pi)c \\
&= \lambda\mu_{hn}\theta[(a_2-1)\Delta W + b_2] - \lambda\mu_{lo}\theta[(1-a_1) - b_1] + \lambda(\mu_{hn}+\mu_{lo})\Delta W \\
&\quad + \rho\pi(1-\pi)\phi(\nu_n - \nu_l) - (1-\pi)c
\end{aligned}$$

Thus,

$$\Delta W = \frac{\rho\pi(1-\pi)(\nu_n-\nu_l)\phi + \lambda\mu_{hn}\theta b_2 + \lambda\mu_{lo}\theta b_1 - (1-\pi)c}{r+\lambda\mu_{hn}\theta(1-a_2) + \lambda\mu_{lo}\theta(1-a_1)}$$

Let

$$\begin{aligned}
F(\rho,\pi,c) &\equiv \frac{\rho\pi(1-\pi)(\nu_n-\nu_l)c}{r+\lambda(\mu_{hn}+\mu_{lo}) + \rho\nu_n\pi + \rho\nu_l(1-\pi)} - (1-\pi)c \\
&= c\pi\Big[\frac{\rho(1-\pi)(\nu_n-\nu_l)}{r+\lambda(\mu_{hn}+\mu_{lo}) + \rho\pi\nu_n + \rho(1-\pi)\nu_l} + 1\Big] - c \\
&= c\pi\Big[\frac{r+\lambda(\mu_{hn}+\mu_{lo}) + \rho\nu_n}{r+\lambda(\mu_{hn}+\mu_{lo}) + \rho\pi\nu_n + \rho(1-\pi)\nu_l}\Big] - c
\end{aligned}$$

and thus

$$\Delta W = \frac{F(\rho,\pi,c) + \lambda\mu_{hn}\theta b_2 + \lambda\mu_{lo}\theta b_1}{r+\lambda\mu_{hn}\theta(1-a_2) + \lambda\mu_{lo}(1-a_1)}$$

B.3.2 Impact of ρ on ΔW and Bid-Ask Spread

Let

$$\omega = \frac{r + \lambda(\mu_{hn} + \mu_{lo})}{\rho}$$

Then, $F(\rho, \pi, c)$ can be written as

$$\tilde{F}(\omega, \pi, c) = c\pi\Big[\frac{\omega + \nu_n}{\omega + \pi\nu_n + (1-\pi)\nu_l}\Big] - c$$

Since $\nu_n > \nu_l$, $\tilde{F}(\omega, \pi, c)$ is decreasing in ω and thus increasing ρ. It is then immediate that $\partial\Delta W/\partial\rho > 0$. From Lemma 2, the bid-ask spread is decreasing in ΔW, and thus the bid-ask spread is decreasing in ρ.

B.3.3 Impact of c on ΔW and Bid-Ask Spread

On the other hand,

$$\begin{aligned}
\frac{\partial F(\rho, \pi, c)}{\partial c} &= \pi\Big[\frac{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\nu_n}{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\pi\nu_n + \rho(1-\pi)\nu_l}\Big] - 1 \\
&= -\frac{(1-\pi)[r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\nu_l]}{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\pi\nu_n + \rho(1-\pi)\nu_l} \\
&< 0
\end{aligned}$$

Therefore, $\partial\Delta W/\partial c < 0$, that is, the dealers' ex-ante expected reservation value ΔW is decreasing in c. Thus, the bid-ask spread is increasing in c.

B.3.4 Interaction between ρ and c

Previously we have

$$\begin{aligned}
\frac{\partial F(\rho, \pi, c)}{\partial c} &= \pi\Big[\frac{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\nu_n}{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho\pi\nu_n + \rho(1-\pi)\nu_l}\Big] - 1 \\
&= \pi\Big[\frac{\omega + \nu_n}{\omega + \pi\nu_n + (1-\pi)\nu_l}\Big] - 1
\end{aligned}$$

which is decreasing in ω and thus increasing in ρ. Hence,

$$\frac{\partial^2 \Delta W}{\partial \rho \partial c} > 0$$

Since ΔW is increasing in ρ, the result means that as c increases, the magnitude of the impact of *rho* on ΔW also increases. On the other hand, since ΔW is decreasing in c, as ρ increases, the magnitude of the impact of c on Δ decreases. From Lemma 2, it is immediate that the magnitude of the effect of ρ on the bid-ask spread is also increasing in c, and the magnitude of the effect of c on the bid-ask spread is decreasing in ρ.

B.4 Proof of Proposition 4

B.4.1 Relationship Between ΔW and Hedging

From Appendix B.3, we have

$$F(\rho, \pi, c) = c\pi \Big[\frac{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \nu_n}{r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \pi \nu_n + \rho(1-\pi)\nu_l} \Big] - c$$

Thus,

$$\frac{\partial F(\rho, \pi, c)}{\partial \pi} = c \frac{[r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \nu_n][r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \nu_l]}{[r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \pi \nu_n + \rho(1-\pi)\nu_l]^2} > 0$$

Thus, $\partial \Delta W / \partial \pi > 0$. Dealers' ex-ante expected reservation value ΔW is increasing in the hedging access π.

B.4.2 Relationship Between Bid-Ask and Hedging

From Lemma 2, we know that the bid-ask spread is linearly decreasing in ΔW. Using the results above, we immediately know that the bid-ask spread is decreasing in π.

B.4.3 Nonlinear Effect of π

From above

$$\frac{\partial F(\rho, \pi, c)}{\partial \pi} = c \frac{[r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \nu_n][r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \nu_l]}{[r + \lambda(\mu_{hn} + \mu_{lo}) + \rho \pi \nu_n + \rho(1-\pi)\nu_l]^2}$$

It is easy to verify that

$$\frac{\partial^2 F(\rho,\pi,c_H,\delta)}{\partial \pi^2} < 0 \quad \text{so} \quad \frac{\partial^2 \Delta W}{\partial \pi^2} < 0$$

Thus, dealers' ex-ante expected reservation value ΔW is strictly concave in the hedging access π. That is, when π is high, ΔW is less sensitive to changes in π, in the sense that the effect of π on ΔW is lower in magnitude. Since the bid-ask spread is linearly decreasing in ΔW, the second derivative of the bid-ask spread with respect to π is positive. That is, when π is high, the effect of π on the bid-ask spread is lower in magnitude.

B.5 Proof of Proposition 5

From Appendix B.4, we have

$$\frac{\partial F(\rho,\pi,c)}{\partial \pi} = c\frac{[r+\lambda(\mu_{hn}+\mu_{lo})+\rho \nu_n][r+\lambda(\mu_{hn}+\mu_{lo})+\rho \nu_l]}{[r+\lambda(\mu_{hn}+\mu_{lo})+\rho\pi\nu_n+\rho(1-\pi)\nu_l]^2}$$

Thus,

$$\frac{\partial^2 F(\rho,\pi,c)}{\partial \pi \partial c} = \frac{[r+\lambda(\mu_{hn}+\mu_{lo})+\rho \nu_n][r+\lambda(\mu_{hn}+\mu_{lo})+\rho \nu_l]}{[r+\lambda(\mu_{hn}+\mu_{lo})+\rho\pi\nu_n+\rho(1-\pi)\nu_l]^2} > 0$$

Thus,

$$\frac{\partial^2 \Delta W}{\partial \pi \partial c} > 0$$

Because $\partial \Delta W/\partial c < 0$, when π increases, the magnitude of the effect of c on ΔW decreases. Thus, the dealers' ex-ante expected reservation value ΔW is less sensitive to changes in c when π is high. Similarly, the bid-ask spread is less sensitive to changes in c when π is high. Conversely, because $\partial \Delta W/\partial \pi > 0$, when c increases, the magnitude of the effect of π on ΔW increases. Thus, when c is high, the dealers' ex-ante expected reservation value ΔW (and the bid-ask spread) is more sensitive to changes in π.

C Additional Data and Empirical Results

C.1 ETF Creation and Redemption Mechanisms

At the end of the day, ETF shares are created or redeemed through the creation and redemption mechanism, facilitated by authorized participants (APs). A novel aspect of the creation and redemption mechanism emphasized by many related papers (e.g., Pan and Zeng (2019), Dannhauser and Hoseinzade (2021)) is the "in-kind" creation and redemption of ETF shares. In an in-kind creation transaction, an AP delivers a basket of underlying bonds (creation basket as specified by the ETF issuer) to the ETF issuer in exchange for shares of the ETF; in an in-kind redemption transaction, an AP delivers ETF shares to the ETF issuer in exchange for a basket of bonds (redemption basket as specified by the ETF issuer). The creation and redemption mechanism is closely related to the APs' ETF arbitrage activities. When the share price of an ETF rises above its NAV, APs have an incentive to engage in the creation process and exchange the creation basket for ETF shares. Conversely, when the ETF share price drops below its NAV, the APs would like to engage in the redemption process and swap shares of the ETF for redemption baskets of underlying bonds.

Not all ETFs have "in-kind" creation and redemption mechanism, as some ETF issuers create or redeem shares exclusively in cash while some others use a hybrid of in-kind and cash creation and redemption mechanisms. Of the 167 ETFs, 61 rely on fully in-kind creation and redemption mechanism while 152 have creation and redemption mechanisms that are either in-kind or hybrid. In contrast to in-kind creation and redemption processes, cash creation and redemption processes do not involve transfer of underlying bonds. Instead, an AP delivers cash to the ETF issuer in exchange for ETF shares in a creation transaction, or delivers ETF shares for cash in a redemption transaction. Thus, ETFs relying on the cash creation and redemption mechanism are akin to traditional open-ended mutual funds. Table C.1 presents descriptive statistics for the corporate bond ETFs that rely on different creation and redemption mechanisms. Fund total assets, shares outstanding and options data are from Bloomberg while secondary trading volume data are from CRSP. The ETFs with only in-kind creation and redemption mechanism tend to be much larger in terms of both total assets and shares outstanding, and more likely to have options available. ETFs with cash-only creation and redemption mechanism tend to have lower primary and secondary volumes, in terms of shares created or redeemed in the primary market or shares traded in the secondary market. These ETFs also tend to trade less frequently in both primary and secondary markets, in terms of shares turned over.

Table C.1: Types of ETF Creation and Redemption Mechanisms

Type	Obs	Median AUM	Shares Out	% with Options	Median Volume
In-Kind	61	581.61	23.00	21.3%	6.70
Hybrid	90	141.56	4.15	12.2%	0.79
Cash	16	82.94	2.63	6.3%	0.44
Total	167	159.21	4.90	15.0%	0.94

Note: The table presents descriptive statistics for corporate bond ETFs with available data on fund total assets (AUM) from Bloomberg as of June 30, 2020. There are three ETF creation and redemption mechanisms - in-kind only, cash only, and a hybrid mechanism involving a mix of the two. The median AUM (in $mn) is the median total assets on June 30, 2020, which are obtained from Bloomberg. Shares outstanding represents the median number of shares outstanding as of June 30, 2020 in millions. The data on shares outstanding are obtained from Bloomberg. % with options represents the percentage of ETFs that have options available. The data on options are also obtained from Bloomberg. Median Secondary means the median (across ETFs) of the average monthly trading volumes between January 1, 2020 and June 30, 2020, in terms of total shares (in millions) traded on the secondary market. Secondary trading volumes are obtained from CRSP.

The ETFs with only in-kind creation and redemption mechanism tend to be much larger in terms of both total assets and shares outstanding, and more likely to have options available. ETFs with cash-only creation and redemption mechanism tend to have lower trading volumes as well. In fact, the 61 ETFs with fully in-kind creation and redemption mechanisms held 86% of the total assets across all of the 167 corporate bond ETFs, as of June 30, 2020. Furthermore, the top 10 largest ETFs in terms of total assets accounted for around 80% of total assets across the corporate bond ETFs, and all of the top 10 largest ETFs rely solely on in-kind creation and redemption process.