Portfolio Trading in Corporate Bond Markets

Jeffrey Meli *†

Barclays, US

Zornitsa Todorova*

Barclays, UK

21st December 2023

Abstract

Portfolio trading, a recent innovation in the corporate bond market, involves trading a basket of bonds as a single piece of risk with a single market-maker. We develop an algorithm to identify portfolio trades in TRACE, and show that portfolio trading increased from 1% of total investment grade corporate bond volumes in 2018 to 7% in 2021. The protocol reduces execution costs by over 40%, with the largest benefits accruing to the least liquid bonds. We link these gains to the ETF ecosystem; portfolios that are more easily priced and hedged using ETFs have lower transaction costs.

JEL Codes: C55, G12, G14

Keywords: portfolio trades, corporate bonds, transaction costs, ETFs

---

We thank Itzhak Ben-David, Adam Kelleher, Arik Ben Dor, Carlo Favero, Claudio Tebaldi, Daniele d'Arienzo, Melissa Prado, Nicholas Hirschey and seminar participants at Columbia Business School, Bocconi University and Nova School of Business and Economics for helpful comments.

† Corresponding author: Jeffrey Meli, Tel: +1 212 412 2127, email: *jeff.meli@barclays.com*

Barclays, 745 7th Ave, New York, NY 10019

1. Introduction

Portfolio trading is a new trading protocol in the corporate bond market, in which an investor bundles a set of individual corporate bonds into one basket and executes the entire basket as a single piece of risk, with one market-maker.[1] In this article we make both methodological and analytical contributions towards a better understanding of this innovation. Our methodological contribution is to construct a comprehensive database of portfolio trades (PTs) going back to their inception in 2018, using insights from a proprietary sample of portfolio trade inquiries received by a large market-maker. Using this dataset, we show that PTs are remarkably effective; they reduce realized transaction costs by more than 40%, with the bulk of the gains accruing to the least liquid bonds. We link these reduced transaction costs to the corporate bond ETF ecosystem. Corporate bond ETFs provide market-makers with real-time price transparency and hedging capability for ETF-like portfolios of corporate bond risk, as well as an outlet for less liquid bonds through the create and redeem (C/R) process. Together, these drive the reduction in PT execution costs.

The link between portfolio trading and the ETF ecosystem is the key insight of this article. Corporate bond liquidity is constrained by the vast number of unique bond (CUSIPs); unlike equities, each bond issuer can have many securities outstanding, each of which has a different maturity and coupon, and possibly different seniority, optionality, and covenants. The large number of CUSIPs makes it difficult to match buyers and sellers, and explains why corporate bond trading is still primarily done via bi-lateral OTC transactions with market makers, rather than via "all-to-all" or other equity-like trading venues. A market-maker who transacts in a particular CUSIP must either hold that position in its inventory or find the other side of the trade. The post-crisis regulations raised the cost of holding inventory, which

[1] This contrasts with the standard request-for-quote (RFQ) protocol, under which investors execute each bond trade individually.

reduced liquidity and made execution more reliant on matching buyers and sellers (see Goldstein and Hotchkiss (2020), among others).

PTs turn this constraint on bond liquidity into a strength. Investors naturally own many (often thousands) of CUSIPs in their portfolios, and there are many aspects of corporate bond portfolio management that can be accomplished, or even require, trading large numbers of bonds, such as adjusting curve, sector, or credit quality exposures. Due to the deep secondary market liquidity of corporate bond ETFs (Meli and Todorova, (2023)), they provide real-time pricing and hedging tools for portfolios of bonds. This is particularly true for portfolios that are highly correlated to the ETFs; one of our key results is that PTs that are "closer" to ETFs get better execution. Our analysis exposes a new channel through which corporate bond ETFs affect their underlying securities. ETFs take a step towards "completing" the corporate bond market, by facilitating more efficient transactions of specific types of risk that would be otherwise difficult to trade.

This analysis requires constructing a database of PTs using TRACE, which is challenging because they were not flagged in the TRACE feed until May, 2023 (the addition of a PT flag reflects the growing importance of this new protocol).[2] To do so, we use a proprietary database of PT inquiries in investment grade (IG) corporate bonds received by a large market-maker. We match these inquiries to TRACE to find a verified set of PTs that actually traded. We use those verified PTs to develop a machine learning clustering algorithm to identify additional portfolio trades that are not already part of our inquiry database.[3] The result is a fulsome dataset of portfolio trades spanning more than 12,000 unique IG PTs and c.1 million bond-PT transactions. We perform a number of validation checks on this dataset

[2] The reporting rule changed on the 15th May 2023, as described in FINRA's *Regulatory Notice 22-12*.

[3] Any one market maker would only receive a subset of PT inquiries, thus the list of verified PTs we identify is necessarily incomplete; the algorithm we develop allows us to overcome this limitation.

to ensure that the algorithm identifies actual portfolio trades. Using this novel database, we show that portfolio trading has grown rapidly, from 1% of volumes in 2018 to 7% of volumes in 2021. This rise mirrors the increase over time in the number of inquiries in our database and demonstrates why the activity level is high enough to justify the addition of a PT flag in TRACE.

By comparing the characteristics of the bonds included in portfolio trades to those of bonds executed in RFQs we conclude that the main motivation for PTs is to improve the ability to trade illiquid bonds. First, PTs tend to have lower liquidity than the trades done in the standard RFQ format. Second, this reduced liquidity is not uniform across the typical portfolio. Instead, portfolios combine some very illiquid bonds with other highly liquid bonds. We interpret this as an attempt to "crowd-source" liquidity in illiquid bonds whereby the overall transaction cost is reduced by bundling them with more easily traded securities. Finally, even in portfolios that appear to target specific market segments (such as the long end of the credit curve) most of the portfolios heavily feature illiquid bonds.

To assess the effectiveness of portfolio trading, we employ a rigorous regression specification, where we compare the difference in the realized transaction cost of portfolio trades and standard RFQs for the same bond on the same day. We control for trade-level characteristics, and include bond-date fixed effects. In aggregate, we find that transacting via a portfolio trade reduces costs by over 40%. The reduction in transaction costs is not uniformly distributed across bonds. The greatest benefit accrues to the least liquid bonds; all else equal, portfolio trading is two to three times more cost effective for illiquid than for liquid bonds. This result holds for a diverse set of liquidity measures such as transaction costs, trade volume, price impact, autocorrelation in returns and bond age.

That PTs reduce transaction costs explains their popularity with investors, but leaves unanswered the question of why they are so effective. Phrased differently, why are market-makers willing to execute them at such low transaction costs? We argue that the benefits of PTs depend critically on the ETF ecosystem. Several statistics are suggestive of this link. First, on average, 60% of the line items included in PTs are owned by the largest IG ETFs, whereas these ETFs only own about 30% of the bonds in the broad IG corporate bond index. Second, at a bond level, the ratio of PT volumes to total volumes is increasing in both ETF ownership and bond illiquidity. For example, for very liquid bonds with low ETF ownership, only about 6% of their trading volume is done via portfolio trades. For illiquid bonds with high ETF ownership, 15% of their volume is done via portfolio trades.

One benefit of ETFs is that they allow market-makers to easily price and hedge ETF-like baskets of corporate bonds; in contrast, ETFs have limited utility for pricing and hedging individual RFQs. We use our regression model to predict the transaction cost of a PT if it had traded as a series of individual RFQs, and define the "PT benefit" as the difference between the actual PT transaction cost and this predicted RFQ cost. Equipped with this measure, we show that the benefit of executing via a PT increases with the "closeness" of the portfolio to ETFs. For example, an interquartile shift in the correlation of portfolio returns to the returns of LQD (a large IG ETF) increases the PT benefit by 12%.

ETFs also routinely transact in bonds they own via the create and redeem process, providing an outlet to offload inventory accumulated from portfolio trades. This is particularly important for illiquid bonds, which might otherwise be difficult or costly to trade, and for which the benefits of transacting via a PT are particularly strong. In cross-sectional regressions, we show that being "right way" vis-à-vis the ETFs, meaning that investors are selling bonds in the ETF create basket and buying bonds in the ETF redeem basket, reduces transaction costs for the least liquid bonds.

Relationship to prior literature

Our analysis contributes to several areas of the existing literature. First, we contribute to the literature that studies how the supply of and demand for corporate bond liquidity has evolved since the global financial crisis. Several papers demonstrate that corporate bond liquidity deteriorated in the aftermath of the crisis (e.g., Dick-Nielsen, Feldhütter, and Lando (2012); (Friewald, Jankowitsch, & Subrahmanyam (2012); (Bessembinder, Jacobsen, Maxwell, & Venkataraman, (2018)). Against this backdrop, a large body of work investigated how the supply of liquidity provided by market-makers has changed with market conditions, regulations and trading protocols (e.g., Goldstein and Hotchkiss (2020), Goldberg and Nozawa (2020) and Carapella and Monnet (2020) among others). For example, as market-makers became less willing to hold inventory, more trades were done on an agented basis (meaning market-makers line up the other side of the trade before executing), which involves a trade-off between transaction costs and immediacy and certainty of execution. Other research has instead focused on how investors adapt to lower liquidity. Jiang, Li, and Wang (2021) demonstrate that open-end corporate bond funds dynamically manage liquidity to meet investor redemptions; Meli and Todorova (2023) show that high yield mutual funds use ETFs to manage liquidity, which results in an aggregate decline in high yield bond liquidity as investors substitute trading in ETFs for trading in the underlying bonds.

Our analysis documents the next stage in the development of new trading protocols and the management of liquidity needs, as investors take advantage of new developments in the market to mitigate the effect of reduced liquidity. PTs allow investors to efficiently trade less liquid securities. This benefits active managers, who can more easily position around their benchmark, and could reduce the liquidity premium in corporate bonds.

We also contribute to the literature on the implications of ETFs for the underlying financial markets. In equities, the literature shows that ETFs have a positive effect on volatility (Ben-David, Franzoni, and Moussawi (2018)), return co-movement (Da and Shive (2018)) and liquidity co-movement (Agarwal et al. (2018)). For bond ETFs, several studies show that ETFs lead to better liquidity (e.g. Holden and Nam (2019), Ye (2019), Marta (2020), Meli and Todorova (2023)) and better price discovery (Choi, Kronlund and Oh (2022)), but could weaken bond price informativeness (Rhodes and Mason (2022)) and increase bond fragility (Dannhauser and Hoseinzade (2022)). Shim and Todorov (2021) document that ETF C/R baskets are fractional and discuss implications for ETF premiums and discounts. Koont et al. (2022) show that basket inclusion generates additional trading activity, which improves the liquidity of the bonds in the ETF baskets.

We identify a novel channel through which ETFs affect underlying markets: they facilitate the pricing and hedging of diversified portfolios of corporate bond risk, particularly when the portfolio is "close" to the ETF. This channel is uniquely suited to the corporate bond market. It allows investors and market-makers to abstract away from the onerous matching problem caused by the large number of CUSIPs, and thus facilitates transactions in individual line items that would otherwise be expensive or impossible to execute. Interestingly, this channel is indirect: PTs rely on the existence and liquidity of ETFs, despite the fact that the investors utilizing the protocol need not ever buy or sell an ETF directly. This also speaks to the limitations of the protocol; it works least well when ETF flows are one-sided and the hedging capabilities are impaired. For example, initial results indicate that PTs did not have lower transaction costs during March 2020, when volatility spiked due to the Covid-19 outbreak, and flows were one-sided (Meli and Todorova (2022b)).

Finally, our paper is most closely related to work by Li et al. (2023), who apply our algorithm to study portfolio trades in TRACE. While the authors also find that portfolio trades improve corporate bond liquidity, they relate the benefits to diversification and not the ETF ecosystem. These differences can be fully explained by sample construction. Li et al.

(2023) include the early years of portfolio trading (2018-2020) in their analysis and use both investment grade and high yield bonds, whereas we base our results on IG portfolio trades executed in 2021. First, our proprietary database of portfolio inquiries shows that 20% of the count of portfolio trades during 2018-2020, compared to less than 10% in 2021 were "custom" portfolios. In a custom portfolio, an investor approaches a market-maker to construct a portfolio which meets a specific investor objective. The line items the market-maker includes in the portfolio typically reflect their inventory and the final portfolio passes several rounds of discussions and negotiations. As technology improved and the PT market matured, in 2021, the majority of PTs were done in-competition – a fast-paced process which requires pricing a basket of bonds in a matter of several minutes and where we see a much bigger role for the intra-day price transparency of ETFs. Finally, the importance of the two linkages to the ETF ecosystem could be different in IG and HY. While the ETF hedging and transparency arguments likely applies equally to HY PTs, the ETF C/R channel as an outlet for illiquid bonds could be less important in HY. Preliminary evidence (Meli and Todorova (2022a)) suggests that HY PTs do not over-index in illiquid bonds, and since all the large HY ETFs are benchmarked against liquid indices, we would expect a more limited role for the ETF C/R mechanism for HY PTs. Studying in detail differences between the IG and HY PT market is an interesting topic for future research.

2. Data and Variables Definitions

2.1 Portfolio Trade vs. RFQ Protocol

In the standard RFQ protocol, an investor requests a quote on a bond for a specific trade size from one or several market-makers, and typically transacts (if at all) with the market-maker that provided the best price. Investors at times make RFQs for a large number of bonds at once.[4] The individual trades that comprise these lists are typically referred to as line items

[4] These lists are known as "BWICs" (bids wanted in competition) or "OWICs" (offers wanted in competition).

and the responses are evaluated on an item-by-item basis; the investor executes each line item individually with the market-maker that provided the best quote, with no expectation that the transactions will be pooled or bundled.

In a portfolio trade investors ask market-makers to price an entire portfolio as a single piece of risk. If the investor agrees to the price, the portfolio trade is then executed in its entirety with a single market-maker. Like with an RFQ, an investor can request a quote on a portfolio from one or several market-makers. Although a single price is agreed to for the entire portfolio, each individual line item is still subject to the TRACE (Trade Reporting and Compliance Engine) reporting rules. The prices of the individual line items reported to TRACE (weighted by their respective notionals) must sum to the quoted price of the portfolio.

There are several reasons why both investors and market-makers ensure that the prices reported to TRACE are accurate at the bond level (i.e., that the portfolio price is correctly apportioned across the individual line items). For example, investors have best execution requirements that apply at the bond, and not at the portfolio, level. An investor who bought a portfolio where some trades were priced too richly would attract scrutiny, regardless of whether others in the same portfolio were priced cheaply. Similarly, market-makers are subject to fair dealing requirements. For example, a market-maker that reported a $500K bond purchase to TRACE at an artificially low price could be accused of excessive mark-up if it then sold the same bond in the same size at the correct price to a different investor. Therefore, the transaction costs of PTs can, in principle, be compared to those of RFQs at the trade level. That said, PTs we not flagged in TRACE until May, 2023, which presents a practical challenge to any such comparison to RFQs.

2.2 Portfolio Trade Inquiry

To identify PTs, our first step is to collect a proprietary dataset of investment grade (IG) investor portfolio inquiries received by the Barclays corporate bond trading desk over the period 1st October 2018- 31st December 2021. Our sample contains all inquiries received by the trading desk, regardless of whether or not Barclays executed the trade. The inquiry dataset spans c. 3,000 investor inquiries that contain c.22,000 unique investment grade bonds. For each inquiry we obtain the date when it was received (but not an intra-day time stamp), and for each line item in the inquiry, the unique bond identifier (CUSIP), trade size (ie. notional), and direction (buy or sell). ***Table I*** gives an example of a typical inquiry.[5]

The number of portfolio trade inquiries in the dataset grew significantly over the sample period, from virtually zero in early 2018 to more than 2,000 inquiries and $175 billion in inquiry volume in 2021 (***Figure 1***). While we believe that Barclays has a large enough market share such that the sample of inquiries is representative in terms of line items, trade sizes, and execution times, the inquiry dataset is not a full accounting of all PT inquiry in the market. No single market-maker has access to the complete set of inquiries because institutional investors balance the potential for price improvement from submitting their prospective portfolio trade to many counterparties against the risk of revealing market-moving information. This motivates the need to construct a more comprehensive database of portfolio trades. Hence, in ***Section 3***, we use the proprietary inquiries database to develop an algorithm which can identify portfolio trades which are not included in our inquiry data.

---

[5] The dataset also contains a flag if the portfolio is "custom" or "in-competition". We discard the custom portfolios, which make up about 10% of the 2021 sample and 20% of the 2018-2020 sample, because they are designed by the market maker at the request of the investor to achieve a specific investment objective (e.g., the investor wants to buy $150 million BBB-rated, 12+ maturity debt). Since it is possible that the line items in these custom inquiries are influenced by the market-maker's existing inventory and risk appetite, they might not be representative of the market.

2.3 Bond Sample and Liquidity Measures

Bond Sample

We obtain transaction-level corporate bond data from the standard version of TRACE, which caps trade sizes at $5 million for IG bonds, for the period 1st October 2018- 31st December 2021.[6] We follow the approach by Dick-Nielsen ((2009), (2014)) to remove double counting, corrections, reversals and cancellations from TRACE. We then augment the cleaned TRACE data with bond-level characteristics from Bloomberg (spread, maturity, time since issuance, numeric rating, amount outstanding, issue size, sector classifications and call types), computed at the beginning of each month. The Bloomberg data covers dollar-denominated bonds belonging to major bond indices (e.g., Bloomberg Investment Grade Corporate Bond Index). We drop bonds with incomplete or missing data. The resulting bond data contains records for 97% of the line items in the portfolio inquiry dataset.

Liquidity Measures

We compute five liquidity metrics at the bond-month level: Liquidity Cost Score (LCS), bond age, Trade Efficiency Score (TES), Price Impact, and Roll's measure.

LCS is a commercially available measure of transaction cost computed using quotes from the Barclays trading desk. It follows the methodology by Konstantinovsky, Yuen Ng, and Phelps (2016). LCS measures the transaction cost for an institutional-size round-lot trade, expressed as a percentage of the bond's price (hence higher LCS signifies lower liquidity). We also use bond age as a proxy for liquidity, the intuition being that bonds are most liquid shortly after issuance, and as bonds age, their liquidity tends to decrease.

[6] Since individual line items in a portfolio trade rarely exceed the cap, working with the standard TRACE data instead of the enhanced version does not have a material impact on our analysis.

The other measures are computed using transaction data from TRACE.[7] TES blends transaction costs and trading volume into a single relative trade score, reflecting both the cost and the flow. To calculate TES, we assign each bond to a monthly LCS quintile and a monthly trading volume decile. Then, the sum of these quantiles (ranging from 2 to 15) is mapped to a TES ranking from 1 to 10, where a lower TES corresponds to better liquidity. We define Amihud's (2002) daily measure of price impact as the volume-weighted absolute daily return. Then, to convert to a monthly frequency, we use the median value of the daily price impact in that month. Finally, in the spirit of Roll (1984), we compute the first-order auto covariance using all transaction level price changes within a given month.

There are advantages and disadvantages of both quotes-based and trade-based liquidity measures (Schestag, Schuster, & Uhrig-Homburg (2016)). The advantage of quotes-based measures is that they are not limited to realized transactions only. The concern is that quotes from a single market-maker on an individual bond could reflect the inventory and risk appetite of that particular market-maker. This could introduce noise but not bias, as any positioning or risk tolerance will average out over time and across bonds. The advantage of liquidity measures computed from TRACE data is that they include trades from all market-makers and are not influenced by a single market-maker's inventory position. On the other hand, the typical concern about TRACE-based liquidity measures is that they are computed from realized transactions only, which may present a distorted picture of liquidity. For example, trading by "appointment" or on an "agented" basis, in which a market-maker only executes a trade after both sides have been identified, incurs lower transaction costs than trading on a "principle" basis, in which a market-maker takes a risk position (Kargar et al.

[7] For other liquidity measures refer to work by (e.g. Puh (2009), Feldhütter (2012), Dick-Nielsen, Feldhütter, and Lando (2012), Corwin and Schultz (2012)), and frequency of trading (also known as zeros) (Lesmond, Ogden, and Trzcinka (1999)).

(2021)). Combining quotes-based and transaction-based measures helps overcome these concerns and paints a more holistic picture of corporate bond liquidity.

2.3 ETF C/R Baskets

We construct daily ETF C/R baskets for LQD, the largest IG ETF. Following the methodology by Shim and Todorov (2021) and Koont et al. (2022), we impute LQD's realized creation and redemption baskets from daily changes in holdings on days with C/R activity. Daily ETF holdings are publicly available and can be downloaded from the iShares website.[8] We identify create (redeem) days as those days on which there was a positive (negative) change in the number of LQD shares. We then use daily changes in the number of bonds held to infer the composition of the average LQD basket on each day.

It is possible that there are redeem baskets on days with net creations and create baskets on days with net redemptions, and that different authorized particpants (APs) negotiate different baskets with an ETF on the same day. Therefore, our imputed baskets are best interpreted as the average net basket for LQD on each given day. We have verified that the average (monthly) correlation between the actual and imputed LQD flows is close to 0.80, which gives us confidence in our methodology (***Figure A2.1***).

3. Constructing a Database of Portfolio Trades

Our first contribution is a methodological one; we use our proprietary dataset of portfolio inquiries to develop a machine learning algorithm to identify a fulsome set of portfolio trades in the TRACE database. The dataset we construct includes more than 12,000 unique IG PTs and c. 1 million bond-PT transactions.

[8] *https://www.ishares.com/us/products/239566/ishares-iboxx-investment-grade-corporate-bond-etf*

3.1 Methodology

In developing our methodology, we seek to balance classification error against the ability to find as many portfolio trades as possible. We proceed in four steps (***Figure 2***).[9]

In the first step, we match the portfolio inquiries to TRACE to identify which inquiries actually traded. Generally, we either don't find the inquiry in TRACE at all or we find it in full or very nearly so. For example, we find 68% of the inquiries in full. This "take-it-or-leave-it" nature of portfolio trades works to our advantage because it allows us to obtain a clean set of traded inquiries, without worrying that the line items we have not been able to match (for whatever reason) could introduce a large degree of noise in our model. We then analyse the matched inquiries and construct the blueprint of the typical portfolio trade in TRACE in terms of the distribution of execution time stamps, number of line items, volumes, average trade sizes etc.

The trades associated with an individual portfolio trade appear as clusters in the TRACE data, with the same or very similar execution time stamps. Hence, in the second step, we run a machine learning clustering algorithm on the TRACE data. This clustering algorithm classifies the TRACE trades into two types of trades: clusters of "candidate" portfolio trades executed within a window of a few seconds, and all other (non-portfolio) trades. The time we allow to elapse between the line items in each candidate portfolio trade is a conservative estimate of the patterns we see in the matched inquiry database.

Third, we re-cluster portfolio trades in order to group together "legs" of the same portfolio trade. This is motivated by the fact that in some cases portfolio trades are reported in TRACE in blocks, separated by a few minutes. This is most common for the different legs

[9] The Data Appendix contains more details on each of these steps. The Python code we used to identify portfolio trades is available upon request. Although we restrict our analysis to the IG market only, the code is designed in a way that allows researchers to construct a HY portfolio trades database as well.

of long-short portfolios; for example, the long leg may be reported at 11:45:10 and the short leg at 11:46:50. If we increased the clustering window in Step 2 the model will correctly identify that both legs belong to the same portfolio, but at the cost of identifying many false positives, which are trades that simply happened to be executed between the times of 11:45:10 and 11:46:50. It is worth noting that if we were only interested in an aggregate estimate of the PT volumes, re-clustering is not necessary. However, analysis of the linkages between PTs and ETFs requires knowing more precisely the composition of each portfolio.

Finally, we apply a series of filters to the data to convert the candidate portfolio clusters into actual portfolio trades. The two most restrictive filters are the exclusion of candidate portfolio trades with fewer than 25 line items and of candidate portfolio trades executed around popular delayed spotting times.

Although c.10% of our inquiries have fewer than 25 line items, we believe many of these are not strictly speaking PTs and are not representative of how PTs are actually priced.[10] This filter is also in part informed by FINRA's definition of portfolio trades, according to which a portfolio trade involves at least 10 unique corporate bonds.[11] We use a stricter definition and apply additional requirements for the minimum trade volume and average trade side in order to ensure that we capture institutional-size transactions and limit classification error.

IG bond volumes in TRACE exhibit sharp daily spikes around known times, which represent delayed treasury spotting.[12] Transactions in IG corporate bonds are often made by counterparties agreeing on a spread to a benchmark Treasury. The actual dollar price of the

---

[10] For example, if an investor mistakenly requests a price for a list of 10 bonds via the PT protocol instead of the RFQ protocol, we would capture the list in the inquiries database, and potentially see the line items subsequently printed in TRACE. However, in reality, this was executed in an RFQ. Alternatively, when investors first use the PT protocol, they typically request PT quotes for a smaller basket of bonds; as they get more comfortable with the process, they increase the size of the PT basket. However, execution quality for a small PT with a limited number of CUSIPS could be very different from the execution of a larger PT.

[11] FINRA's *Regulatory Notice 22-12*

[12] We estimate that in 2021 between 7% - 12% of IG trade volume was printed in TRACE in the 5-minute interval around popular spotting times (i.e. 15:00, 15:30, 16:00).

trade is computed at a later point in time using the previously agreed upon spread. For example, the counterparties agree to a spread at 13:30, but the price is calculated and reported to TRACE at 15:00. Delayed spotting allows investors and market makers to concentrate (and net) their Treasury hedging, instead of having to do multiple hedges throughout the day. While we know from our inquiries database that some portfolio trades are reported around spotting times, the sheer amount of trades that are concentrated around these times makes it impossible to accurately separate portfolio trades from delayed spot RFQ trades. To ensure that our comparison between PTs and RFQs is always meaningful, we also drop RFQ trades executed around spotting times from our regression sample.[13]

3.2 Summary Statistics and Algorithm Validation

The resulting portfolio trades database that we construct contains 12,107 unique IG portfolio trades and c. 1 million bond-portfolio observations, totalling $696 billion of executed bond volumes (Panel A, ***Table II***).

Panel B in ***Table II*** shows how the portfolio trades identified by our algorithm have evolved over time. Portfolio trading activity increased sharply, both in terms of the number of portfolio trades as well as the total dollar volume. We identified 1,950 unique PTs in 2018 and 2.5 times more in 2021 alone (4,914). Portfolio trading volume increased from $81 billion in 2018 to $311 billion in 2021. In percentage terms, the proportion of total trading volume that occurred in the form of PTs in 2021 was close to 7%, off a base of c.1% in 2018. This rapid growth demonstrates that the protocol has been quickly adopted by a large number of market participants, and justifies the requirement to add a PT flag to TRACE trades starting in May, 2023.

---

[13] We do not drop trades that occur in common spotting times from the denominator when we compute the proportion of volumes that occurs in PT form. Therefore, our estimates of the proportion of TRACE volumes that occur in PT form are necessarily conservative.

Despite our use of filters, the concern remains that due to its enormous size, TRACE contains many standard RFQ trades that are clustered by chance in ways that cause us to mischaracterize them as portfolio trades. However, were that to be the case, we should find a consistent flow of PTs in the TRACE data. Instead, the growth of the PT market as identified by our algorithm closely conforms to the growth in the volume of investor inquires, providing a validation check for our machine learning approach.

In ***Table III*** we check how well the algorithm identifies the portfolios in our Barclays inquiries database. We do this check for two sets of inquiries: in-sample (executed in 2021) and out-of-sample (executed in H1 2022). For any given inquiry, the true positives rate is calculated as the number of line items the algorithm identified divided by the total number of line items in that inquiry. The false positives rate is calculated as the number of incorrectly identified line items divided by the total number of line items the algorithm found. In-sample, the median true positives rate is 97% and the false positives rate is 3%; out-of-sample the median true positives rate is 100% and the false positives rate is 7%. These results give us confidence that the algorithm is successful at identifying actual portfolio trades in TRACE.

In ***Table IV***, we compute portfolio-level summary statistics along two dimensions: portfolio construction characteristics (Panel A: number of line items, volume, line item weights and sectors) and volume-weighted bond characteristics (Panel B: liquidity measured by *LCS*, maturity and bond age). We include statistics for both the full set of portfolios identified by our algorithm and the set of actual investor portfolio inquiries. The empirical distribution of the portfolio trades identified by the algorithm closely matches the distribution of investor inquiries in each of these key aspects, which suggests our algorithm is not mischaracterizing groups of RFQ trades as PTs. The average portfolio trade contains c.100 line items and $70 million worth of notional, approximately equally-split between the bonds in that portfolio. Portfolio trades are well-diversified and, on average span bonds from 12 different sectors. The average portfolio trade costs 0.84% to transact and is comprised of bonds with remaining maturity of about 10 years, issued 2.5 years ago.[14]

Finally, we perform our analysis of transaction costs using both the full dataset of PTs and the narrower set of PTs from the inquiry database, and find similar results, which is again supportive of our algorithm. We prefer the analysis using the full database of portfolio trades because the larger sample size allows us to employ a more rigorous econometric specification and to explore in greater detail the cross-sectional heterogeneity in the data.

4. Characteristics of Portfolio Trades

4.1 Crowd-sourcing Liquidity via Portfolio Trades

A first distinguishing feature of PTs is that they are more concentrated in illiquid bonds. We compare PTs identified by our algorithm to RFQs (defined as those trades not identified as PTs) on volume-weighted liquidity (LCS), maturity, and bond age (for RFQs see the last row *"TRACE ex PT"* of ***Table IV***). Along maturity and age, PTs are quite similar to RFQs. The main difference between them appears to be liquidity: the bonds traded in portfolio trades are substantially less liquid than the trades done using RFQs. The average LCS of the line items in PTs is 0.84% compared to 0.69% for RFQs (higher LCS implies lower liquidity). Further, the distribution of portfolio LCS reveals that this is not driven by a few very illiquid portfolio trades. More than 50% of the portfolios, both by count and by volume, are less liquid than the average bonds traded in RFQs.

In ***Figure 4*** we expand on this result by showing how the aggregate distribution of IG portfolio volumes varies by LCS quintile. As a reference point, we also overlay on the same chart the distribution of volumes for bonds in the Bloomberg US IG Corporate Bond Index

[14] As a further robustness check, in ***Figure 3*** we also overlay the percentage distribution of portfolio volumes by sector for the set of portfolio trades identified by the ML algorithm and the set of investor inquiries, and find no material differences in the sectoral distribution of volumes.

(BBG IG). The BBG IG Index contains about seven thousand bonds from a diverse set of issuers and measures the performance of the investment grade, fixed-rate, taxable US corporate bond market. The index is not skewed towards liquid bonds and is widely considered to be representative of the IG corporate bond market. ***Figure 4*** shows that PT volumes are shifted towards the less liquid quintiles. Compared to the BBG IG Index, there is 6% less portfolio volume in the first two most liquid quintiles, the majority of which then sits in the 4th LCS quintile.

Further, we find that investors construct these illiquid portfolios is a way that "crowd-sources" liquidity for the illiquid CUSIPs. To demonstrate, for each portfolio we compute the percentage of volume contained in the two most liquid quintiles of LCS and in ***Figure 5*** plot the portfolio-level distribution of this percentage separately for liquid and illiquid portfolios. Liquid (Illiquid) PTs have lower (higher) trade volume-weighted LCS than the trade volume-weighted LCS of the bonds in the BBG IG Index. The boxplot shows that even amongst the illiquid PTs, very few PTs contain only illiquid bonds (the median percentage of liquid bonds in illiquid PTs is 37%). In other words, these portfolios appear to be designed such that the more liquid bonds cross-subsidize the less liquid bonds, resulting in an overall portfolio LCS that is closer to the index than if these illiquid bonds were traded individually.

We also consider other possible portfolio construction strategies, including implementing a specific investment objective aligned to maturity, sector, or credit rating (***Table A2.1***). We compute maturity, sector and rating-based Herfindahl scores (HHI), summing the squared percentages of trade volume for each individual portfolio, and compare these scores to the respective HHI score of the Bloomberg IG Corporate Bond Index. We classify portfolio trades into a maturity/sector/rating strategy if the HHI of the portfolio is at least 50% higher than the HHI of the Index. Among these other dimensions, a maturity-type strategy is the most common (35% of portfolios) and is typically focused on longer-dated bonds. However, 80% of these maturity-type portfolios can also be classified as illiquid. Even for portfolios tailored to a specific part of the market, liquidity remains a motivating factor.

4.2 Significant Overlap with the ETFs

A second distinguishing feature of PTs is the high degree of overlap between the line items in PTs and the bonds owned by ETFs. In ***Figure 6*** we plot the overlap with the largest IG ETF (over our sample period), the iShares iBoxx Investment Grade ETF (ticker: LQD). On average, 60% of the bonds in IG portfolio trades are owned by LQD. To put this in perspective, LQD owns about 30% of the bonds in the Bloomberg IG Corporate Bond Index. Over 90% of the portfolios in our sample have an overlap in excess of 30%, which shows that portfolio trades are significantly more concentrated in ETF bonds.

LQD was the first IG corporate bond ETF, and owing to its large size (and liquidity) we believe is representative of the IG ETF market. However, over the last decade, the market has grown tremendously and a large number of new ETFs have been established, each tracking a different segment of the IG corporate bond universe. To get a more complete picture of the ETF market and its potential interactions with the PT market, we also analyse the monthly holdings of all IG ETFs contained in the CRSP Mutual Funds Database. For each bond-month in our sample, we compute the percentage of the amount outstanding that is owned by IG ETFs. We then perform a double-sort by liquidity (based on LCS) and ETF ownership quintiles and compute the average share of PT volumes in each bucket. The proportion of total volumes traded via PTs increases as liquidity declines and as ETF ownership increases (***Table V***). For example, for very liquid bonds with low ETF ownership, only about 6% of their trading volume is done via portfolio trades. For the least liquid bonds with high ETF ownership, 15% of their volume is done via portfolio trades.

4.3 Two Hypotheses on the Use of Portfolio Trades

Our analysis suggests that portfolio trades are more concentrated in illiquid bonds and in bonds with higher ETF ownership. These two defining features of portfolio trades are the basis of our two hypotheses, which we will empirically test in the following sections:

H1: Portfolio trades reduce transaction costs, particularly for illiquid bonds

H2: The success of portfolio trades is driven by linkages to the ETF ecosystem

5. Transaction Costs: PTs vs RFQs

5.1 Econometric Model and Identification

We restrict our analysis of transaction costs to data from 2021 in order to study portfolio trading in a more mature stage of its development. Both the number and volume of PTs in 2021 nearly equal the aggregate PT activity for all other years combined. We believe that the earlier data reflected instances when both investors and market-makers were getting familiar with the protocol, and may not be representative of current PTs.

Following the literature (e.g., Bessembinder (2003); Collin-Dufresne, Junge, & Trolle (2020); Hagströmer (2021)) [15], we measure transaction costs of trade *i* in bond *j* on day *t* by the effective half-spread (*EHS*):

$$EHS_{i,j,t} = D_{i,j,t}(P_{i,j,t} - M_{j,t})/M_{j,t}$$

where $D_{i,j,t}$ is an indicator variable that equals one for investor buy trades and negative one for investors sell trades, $P_{i,j,t}$ is the price at which the trade is executed and $M_{j,t}$ is the end-of-day mid-price as quoted by Bloomberg. *EHS* is an indication of how far traded prices are

[15] Using quotes data on S&P 500 stocks, Hagströmer (2021) shows that the level of the effective bid-ask spread measured relative to the mid-price could overstate the true bid-ask spread. The paper derives conditions under which EHS is not biased and proposes new estimators. While Hagströmer's result may apply to corporate bonds, it does not compromise our results since our main interest is the estimate of the difference in execution costs between portfolio and non-portfolio trades, and not the level, which should cancel any bias.

from the mid-price; smaller *EHS* implies lower transaction costs realized by investors.[16] We then multiply *EHS* by 10,000 to compute transaction costs in basis points (bp).

To address potential selection bias in the use of PTs (e.g., that PTs are more heavily utilized for certain bonds or in certain market environments) we compare the transaction costs for portfolio trades and non-portfolio trades in a formal regression model at the transaction level, with controls for trade-level characteristics and bond-date fixed effects:

$$EHS_{i,j,t} = \beta_1 Portfolio\ Trade_{i,j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t} \quad (\textit{\textbf{Model 1}})$$

where $Portfolio\ Trade_{i,j,t}$ is a dummy variable equal to one when transaction *i* in bond *j* on date *t* is part of a portfolio trade. In our baseline analysis, $Portfolio\ Trade_{i,j,t}$ is defined on the sample of PTs identified by our clustering algorithm. Since the specification includes bond-date fixed effects ($\lambda_{j,t}$), the only other variables we need to control for are at the transaction level, collected in the vector ($Z_{i,j,t}$). Given that previous literature (e.g. O'Hara and Zhou (2021); Choi, Huh and Seunghun Shin (2023)) has found that larger trades incur lower transactions costs, we include a dummy variable equal to one if the notional traded in transaction *i* is greater than $5 million ($Block\ Trade_{i,j,t}$).[17] We also include the lagged *EHS* ($Lag\ EHS_{i,j,t}$) to account for any momentum effects in bond-level transaction costs. Further, we control for any noise introduced by computing *EHS* using end-of-day mid prices by including a dummy variable equal to one for trades executed before 13:00 EST ($Morning\ Trade_{i,j,t}$). Finally, we include two other dummy variables $Sell\ Pressure_{i,j,t}$ and $Buy\ Pressure_{i,j,t}$ which equal to one if the five trades preceding trade *i* are all sells or buys

[16] An alternative way to define the *EHS* is to use the last price in the inter-dealer market as a reference point instead of the end-of-day mid-price (e.g. as in O'Hara and Zhou (2021)). However, same-day inter-dealer transactions do not exist for all bonds in the sample, particularly the less liquid bonds.

[17] We use a dummy instead of a continuous measure of quantity traded because we use the standard version of TRACE, where volumes for IG bonds greater than $5 million are capped.

respectively. Finally, we cluster standard errors both at the bond and date levels to account for correlation over time within a given bond and across bonds on a given date.

The main coefficient of interest in ***Model 1*** is β_1, which gives the difference between the *EHS* of PTs and RFQs. If portfolio trading is cost-effective, we expect $\beta_1 < 0$. Having included bond-date fixed effects, identification of the estimates comes from variation in the transaction costs of those bonds which have trades in both protocols on a given day. We see 21% of the bond-date observations in our sample in both protocols, which is a meaningful portion and supports the empirical validity of our results (***Figure 6***).

5.2 Portfolio Trades Reduce Transaction Costs

We find that PTs are substantially more cost-effective than the standard RFQ protocol ($\beta_1 < 0$) (column (1) ***Table VI***). All else equal, the average transaction cost of a line item in a portfolio trade is 6.0bp cheaper than the same trade in RFQ form. Given an average *EHS* of 14.5bp for RFQs, the effect translates into a 41.7% reduction in transaction costs.

In column (4) of ***Table VI***, we re-estimate ***Model 1*** but define $Portfolio\ Trade_{i,j,t}$ using the sample of PT inquiries. Both the magnitude of the coefficients, their statistical significance and the percentage improvement in transaction costs over the standard RFQ protocol remain unchanged. This is an important robustness check, which re-emphasizes that our algorithm identifies actual PTs.

We perform two other robustness checks, using both definitions of $Portfolio\ Trade_{i,j,t}$. First, we bucket trades into \$0.5 million size buckets and include bond-date-size fixed effects; second, we distinguish between buy and sell trades and include bond-date-size-direction fixed effects (columns (2-3) and column (5-6) ***Table VI***). In these regressions, we are effectively comparing the transaction costs for the same bond, on the same day, in the same size, and in the same direction across the two protocols. The magnitude of the effect remains little changed, and we obtain virtually the same estimate of β_1 regardless of whether we use PTs identified by our algorithm or PTs that are part of our original inquiry dataset. While these specifications are extremely rigorous, it is also worth pointing out that both the econometric identification and the economic validity of the effect become progressively more difficult as the sample of bonds with trades fulfilling these criteria shrinks.[18] For these reasons, we use the model with bond-date fixed effects as our baseline.

5.3 The Cost Benefits Are Strongest for Illiquid Bonds

To examine how the benefit of PTs varies across bonds, we augment ***Model 1*** by interacting the $Portfolio\ Trade_{i,j,t}$ dummy with an illiquidity term[19]:

$$EHS_{i,j,t} = \beta_1 Portfolio\ Trade_{i,j,t} + \beta_2 Portfolio\ Trade_{i,j,t} \times Illiq\ Dummy_{j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t} \quad (\textbf{\textit{Model 2}})$$

where $Illiq\ Dummy_{j,t}$ is a dummy variable that equals one for bonds which belong to the most illiquid quintile of the distribution. We measure illiquidity as one of: *LCS, TES, Bond Age, Price impact* or *Roll's measure* for bond *j* on date *t*. If portfolio trading reduces transaction cost to a greater extent for illiquid bonds, we would find both $\beta_1 < 0$ and $\beta_2 < 0$.

In this specification, we compare two differences: first, the difference in transaction costs when a bond is traded in a portfolio and when it is traded via an RFQ, and, second, the difference in transaction costs of an illiquid bond and a liquid bond. Hence, estimating β_2 relies on an additional source of variation compared to our baseline specification. We not only exploit variation in the transaction cost of a bond depending on the trade protocol, but also use cross-

---

[18] We have also estimated specifications where instead of bond-date fixed effects, we (1) include bond and date fixed effects and (2), drop the fixed effects altogether and saturate the model with a comprehensive set of bond-level and date-level controls. The magnitude of the estimates from these regressions tend to be slightly higher than our baseline. While it is relatively easy to control for time-varying features of bonds (e.g. maturity, rating etc.), modelling time-invariant features (which would be captured by the bond fixed effects) tends to be more difficult due to the complex structure of these securities, which is why we prefer using bond-date fixed effects.

[19] Any variation in $Illiq_{j,t}$ will be subsumed the bond-date fixed effects, which is why we don't need to include the main effect in ***Model 2.***

sectional variation in the liquidity profiles of the bonds we observe on any given day. ***Figure 6*** demonstrates that the former holds. The distribution of portfolio volumes by LCS quintiles in ***Figure 4*** demonstrates the latter; portfolio trading is not exclusively confined to the very liquid or the very illiquid bonds only, but occurs across the entire spectrum of liquidity.

Table VII contains the results of ***Model 2***. The reduction in transaction costs for portfolio trades increases as liquidity declines ($\beta_2 < 0$), regardless of the measure of illiquidity we use For example, it is 2.8bp ($= \beta_1$) cheaper to execute a liquid bond (measured by LCS) in a portfolio trade and 8bp cheaper for an illiquid bond (= $\beta_1 + \beta_2$). Evaluated at the mean of *EHS* for RFQ trades, this translates to 19.3% reduction in transaction costs for liquid bonds and 55.2% for illiquid bonds.

6. Relationship to the ETF Ecosystem

The extent to which portfolio trading reduces execution costs, particularly for illiquid bonds, raises the important question why this new trading protocol works so well. We identify two linkages to the ETF ecosystem that drive the reduced execution costs of PTs. First, ETFs provide market-makers an intra-day hedging and pricing tool for transactions in ETF-like portfolios of corporate credit risk. Second, market-makers leverage the ETF create and redeem process (C/R) to offload the inventory of illiquid bonds which accumulates as a result of portfolio trading and/or to source illiquid bonds sold via PTs.

To demonstrate these linkages, we estimate the benefits of trading via PTs at the portfolio level. To do so, we need an estimate of the counterfactual: where would the portfolio have traded as a series of RFQs? The first input in this estimate is the predicted half-spread, $\widehat{EHS_{i,j,t}}$, for each line item in each portfolio trade. We compute this using our baseline model (***Model 1),*** setting the portfolio trade dummy in that regression equal to zero. We then aggregate $\widehat{EHS_{i,j,t}}$ across all the line items in the PT to arrive at a value-weighted portfolio-level measure, $\widehat{EHS_p}$, where the weights are given by the notional of each line item in the PT. The difference between the actual and the predicted effective half spread (i.e., $EHS_p - \widehat{EHS_p}$) measures the benefit of trading the entire portfolio via a PT, where a lower value implies a greater cost improvement relative to the standard protocol. We use this estimate of the PT benefit to explore the portfolio characteristics that impact execution, and thus determine if and how the improvements in execution are linked to the ETF ecosystem. Specifically, we run a series of cross-sectional regressions of $EHS_p - \widehat{EHS_p}$ on portfolio-level characteristics:

$$EHS_p - \widehat{EHS_p} = \alpha + \beta_k Portfolio\ Characteristic_{k,p} + \epsilon_p \quad (\textbf{\textit{Model 3}})$$

to determine which portfolio characteristics are important for PT execution.

6.1 ETFs and PT Hedging and Pricing

One benefit of ETFs vis-à-vis open-end mutual funds is that they trade actively in the secondary market. Intra-day trading provides both price transparency and hedging tools that are far more applicable to portfolios of corporate bonds than to individual bonds. The intra-day price of an ETF allows market-makers to incorporate real-time information into their PT pricing. It also allows investors to compare the price of the portfolio to the price of the ETF to evaluate execution quality. In contrast, the pricing of a single security will reflect mostly idiosyncratic risks and security-specific supply and demand. Similarly, a diversified basket of corporate bonds can be hedged using an ETF, whereas hedging a single name position with an offsetting position in an ETF incurs substantial basis risk.

If the pricing transparency and hedging flexibility afforded by ETFs allows market-makers to reduce PT transaction costs, then PTs for which these tools are more effective should be more cost effective. To show this, we compute three measures of "ETF hedgeability". Our first measure is *Corr LQD*, computed as the correlation between the daily returns of the value-weighted portfolio and LQD during the 30 days prior to the execution of the portfolio trade. With an average correlation of 0.85, LQD appears to be a good hedge for many of the portfolios in our sample (***Table VIII***). For those PTs that are not highly correlated to LQD, there could be a better hedge among some of the other large ETFs. For example, portfolios that are concentrated in bonds with either long or short remaining maturity could be hedged with Vanguard Long-term Corporate Bond ETF (ticker: VCLT, tracking bonds with more than 10 years of maturity) or Vanguard Short-Term Corporate Bond ETF (ticker: VCSH, tracking bonds with between one and five years of maturity). Therefore, we compute a second measure of correlation with ETFs, *Max Corr Top 10 ETFs*, equal to the highest correlation of portfolio returns with the returns of the top 10 IG ETFs[20].

Finally, we also calculate the average ETF ownership of all the bonds in a given portfolio (*% ETF Ownership*). Higher ETF ownership at the portfolio-level means that a good hedging instrument likely exists, even for portfolios which are concentrated in a particular dimension. While both *Corr LQD* and *Max Corr Top 10 ETFs* are direct tests of the ability to hedge a portfolio with an ETF, *% ETF Ownership* is an indirect test, relying on the logic that the higher the overall ETF ownership of the portfolio is, the higher the chance is that a good hedge exists among the broader universe of IG ETFs.

Regardless of the measure we use, we find that portfolios which are which are more effectively hedged using ETFs incur lower transaction costs (columns (1)-(3) in ***Table IX***). For example, increasing the correlation of portfolio returns with LQD returns by the interquartile range improves PT execution by 12%; similarly, an interquartile shift in ETF ownership results in 15% better execution. These results indicate that the benefits of ETFs apply more to PTs than to individual RFQ trades in the same bonds.

[20] Tickers used are LQD, VCIT, VCSH, IGIB, IGSB, USIG, SPSB, SPIB, VCLT and SCHI.

6.2 Illiquid Bonds and the ETF C/R Process

Another benefit of the high ETF ownership is that market-makers can use the ETF C/R process to offset positions acquired via PTs. Market-makers can deliver the bonds purchased and create ETF shares, recycling the risk accumulated as a result of portfolio trading (or the reverse for bonds sold). Bonds that are heavily owned by IG ETFs have a higher probability of being included in the daily ETF create or redeem baskets than bonds with low ETF ownership.

We define *% LQD C/R* as the percentage of the line items in a given portfolio that are in the imputed create or redeem basket of LQD during the five business days after the portfolio was executed. On average, 36.5% of PT line items are in the weekly C/R baskets (***Table VIII***). Shim and Todorov (2021) show that (differently to equity ETFs) the C/R baskets of fixed income ETFs contain only a fraction of the total number of securities held in the ETF portfolio. We confirm that the weekly LQD C/R baskets contain on average 40% of the total ETF holdings. Given that on average 60% of the bonds in a portfolio trade overlap with the holdings of LQD, we would expect an average overlap of 24% (= 60%*40%) between the portfolios and the C/R baskets, substantially below the actual level. In other words, PTs over-index for bond that are included in C/R baskets, even after adjusting for the fact that they over-index for bonds owned by ETFs.

More importantly, we find that a higher overlap with the ETF C/R process significantly reduces transaction costs for portfolio trades (column (4) ***Table IX***). An interquartile shift in the overlap with the ETF C/R baskets reduces the transaction cost of the portfolio by about 15% compared to the average PT benefit. This result is based on the future (realized) C/R baskets, and thus implies that either market-makers can predict these baskets, or that the basket composition is endogenous in a way that is related to the inventory that market-makers accumulate via the PT process (or both).

We can also look at just the overlap with the C/R basket on the day of the PT. We define a trade in an individual bond that an investor sells to a market-maker as "right way" for the LQD when the bond is part of the create basket that day, which would allow market-makers to deliver it to the LQD and create shares. Conversely, an investor buy trade is "right way" when the bond is part of the redeem basket that day. To aggregate at the portfolio level, we compute *% Rightway LQD C/R* as the percentage of "right way" trades in a given portfolio. Consistent with our intuition, we find that portfolios which are better aligned with the ETF process incur lower transaction costs (column (5) ***Table IX***). An interquartile shift in "right way" equates to about a 9% improvement in transaction costs.

However, the benefit of being "right way" is far stronger for illiquid bonds. To demonstrate, we divide each portfolio into liquid and illiquid buckets (using the median of the LCS sample), and define two measures, *% Liquid Rightway for the ETF C/R* and *%Illiquid Rightway for the ETF C/R*, which give the percentage of total liquid and illiquid bonds respectively that are "right way" for the ETF. The magnitude of the "right way" effect is small and statistically insignificant for the liquid bonds (column (6) ***Table IX***): these bonds are relatively easy to trade and so PT execution depends less on ETF activity. In contrast, the benefit is both larger and statistically significant for illiquid bonds. These are generally difficult to trade and thus PT execution benefits substantially from alignment with ETFs. For example, an interquartile shift in the proportion of illiquid bonds that are "right way" for the C/R process improves PT execution by about 8%. Recall from the prior section that the benefits of PTs are most concentrated in illiquid bonds. It is precisely for these illiquid bonds that the effectiveness of PTs is the most sensitive to alignment with the direction of ETF C/R activity: being "right way" results in substantially lower execution costs. Combining these two pieces of evidence, we conclude that PTs benefit from being aligned to C/R activity to a far greater extent than RFQs.

6.3 Alternative Explanations

Our focus is the difference in execution costs between PT and RFQ trades, and not the level. In fact, the time-series correlation between the execution costs of these two protocols is 0.80, which suggests the presence of common factors the costs for both protocols. Market-makers' bid-offer spreads reflect a number of components, which have been well documented in the literature: inventory and hedging costs (e.g. Goldstein and Hotchkiss (2020)), price transparency (e.g. Edwards, Harris, & Piwowar (2007)), and volatility among others. However, the cost benefit of portfolio trading cannot be attributed to these common factors, since by taking the difference in *EHS*, their effect will cancel out. Based on the evidence we provide, we believe that ETFs provide the most likely explanation why portfolio trades work so well. Nonetheless, we have also considered three alternative explanations: diversification, competition for market share and clients "swapping" portfolios. We find no evidence supporting these theories.

A hallmark of PTs is that they contain a large number of line items with a small size and approximately equal weight (***Table IV***). A well-diversified portfolio likely correlates with the ETF, and hence, it is possible that our measures of ETF closeness in fact reflect diversification. To address this, we define *Diversification* as one minus the ratio of the standard deviation of PT returns in the 30 days before the PT was executed and the weighted average of the standard deviation of the returns of the individual line items over the same period:

$$Diversification = 1 - \frac{PT\ Std}{Wtg\ Std\ of\ Line\ Items}$$

By construction, *Diversification* takes values between zero and one. A value of zero means that the PT is not diversified at all and executing either the portfolio or the individual trades carries the same amount of risk for the market-maker; a value of one on the other hand means that executing in portfolio format achieves perfect diversification.

To test this alternative hypothesis, we regress $EHS - \widehat{EHS}$ on three measures of ETF closeness (*Corr LQD, % ETF Ownership* and *% Rightway LQD C/R*) and control for *Diversification* in the same specification (***Table X***). The statistical significance and magnitude of the ETF variables remain unchanged relative to the estimates in ***Table IX*** (Panel A), which provides strong evidence that the linkages to the ETF ecosystem we uncover are not proxies for diversification. The coefficient on *Diversification* is negative and statistically significant, but its economic importance is half that of the ETF variables (Panel B). For example, an interquartile shift in *Diversification* improves execution by 11% compared to 25% for *% ETF Ownership*. Further, this could in fact be a proxy of its own, for different linkages to the ETF ecosystem. Taken together, the results in ***Table X*** support the conclusion that the ETF ecosystem is the main driver of the benefits of portfolio trading.

Market-makers might use portfolio trades to gain market share, and the competition to win these trades could drive the bid-offer they charge. According to this narrative, the reduction in execution costs could either reflect an increased motivation to win the trades, or it could be linked to some information that market-makers obtain through executing these trades that they would not obtain by unsuccessfully bidding/offering on the same portfolio. To test this theory, in ***Figure 8*** we plot the time series of the difference between the average *EHS* of transactions executed via the PT and RFQ protocol over the period Jan 1st 2020 – December 31st 2021. We also show PT volumes as a percentage of total TRACE IG volumes. The reduction in *EHS* associated with PTs remained stable over a period when the use of portfolio trading rose dramatically. This is the opposite of what we would expect if marker-markers were simply “buying” market share, as they would have increased the benefit of portfolio trading to entice more participation.

Finally, it is possible that investors use PTs to "swap" entire portfolios with other investors, who want to trade the same bonds but in the opposite direction. In this scenario, a market-maker would act as an agent, lining up both sides of the PT and charging a considerably lower bid-offer than if the market-maker had acted as a principal and priced the bonds out of inventory. If this were common, we would find examples of offsetting PT trades in our database. However, we find that this occurs in less than 0.5% of the PTs in our sample (***Table XI***), which speaks strongly against the theory.

7. Discussion

In this article we introduce the concept of portfolio trading, the latest innovation in the corporate bond market, which involves trading a basket of bonds as a single piece of risk, and transacting the entire basket with one single market-maker. Using a proprietary dataset of portfolio inquiries, we develop an algorithm to identify corporate bond portfolio trades in TRACE. We show that investors typically use this new trade protocol to transact in illiquid bonds and that portfolio trades reduce transaction costs by more than 40% compared to trades using the standard RFQ protocol. We also demonstrate that linkages to the ETF ecosystem provide a transparent intra-day tool to price and hedge PTs and allow market-makers to both offload the inventory of illiquid bonds which accumulates as a result of portfolio trading.

Our work opens a broad set of avenues for future research. Jiang, Li and Wang (2021) show that when faced with significant redemptions investors typically follow a pecking order of liquidity, selling the most liquid assets first. Meli and Todorova (2023) show that institutional investors use corporate bond ETFs to manage flows-driven liquidity, thus increasingly substituting bond trade volumes with ETFs. In this article, we provide evidence that compared to the standard RFQ volumes, PTs facilitate much more efficient trading in less liquid bonds. Together, these findings suggest greater capacity to manage trading needs at lower cost. The natural question to ask is how the liquidity risk premium has responded to the change in the demand for liquidity as investors have adopted both new products and new ways to manage liquidity and transact in illiquid bonds.

It will be interesting to investigate if and how our results translate to the HY PT market. HY ETFs tend to be even more focused on the liquid spectrum of bonds than IG ETFs, and HY ETFs have a correspondingly lower demand for illiquid HY bonds than do IG ETFs. Another open question is if the rise in PT volumes itself generates spill-overs. The inclusion of less liquid bonds in PTs may give market-makers more comfort providing liquidity in that part of the market, even away from PTs, thereby causing liquidity to snowball.

Finally, we are interested in how transaction costs vary with market conditions. Meli and Todorova (2022b) provide some preliminary evidence that periods when trade volumes are singularly one-sided (investors are either heavily net buying or heavily net selling) coincide with periods when the ETF C/R mechanism is also one-sided, but in a way that most of the trades are “wrong way” for the ETFs. This could potentially limit the ability to offload risk via the C/R mechanism and, thus, also could limit the cost effectiveness of portfolio trades. Relatedly, periods of market distress also typically correlate heavily with periods when the volatility of the ETF bid/offer sharply increases, translating into higher hedging costs of PTs. Market-makers in their turn are likely to pass on these costs to investors, which could result in higher transaction costs of PTs.

List of Figures

Figure 1: Client Inquiries

The figure shows the growth in the number and $ volume of investor portfolio inquiries received by Barclays trading desk.

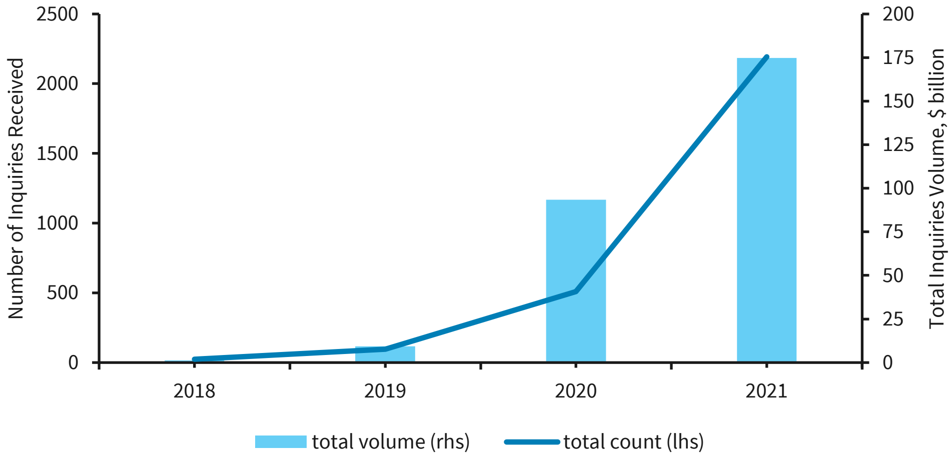

Figure 2: Flowchart of the Methodology Process

The figure shows the steps we undertook to construct the dataset of TRACE portfolio trades.

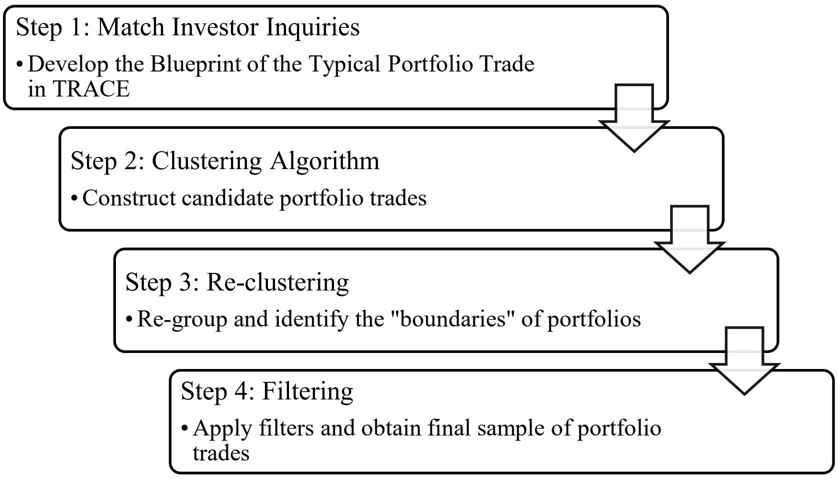

Figure 3: Algorithm Validation – Distribution of Portfolio Trade Volumes by Sector

The figure overlays the distribution of volumes by sector for the TRACE portfolios identified by the ML algorithm and the investor inquiries. Data based on portfolio trades executed during the period 1st January 2021 – 31st December 2021.

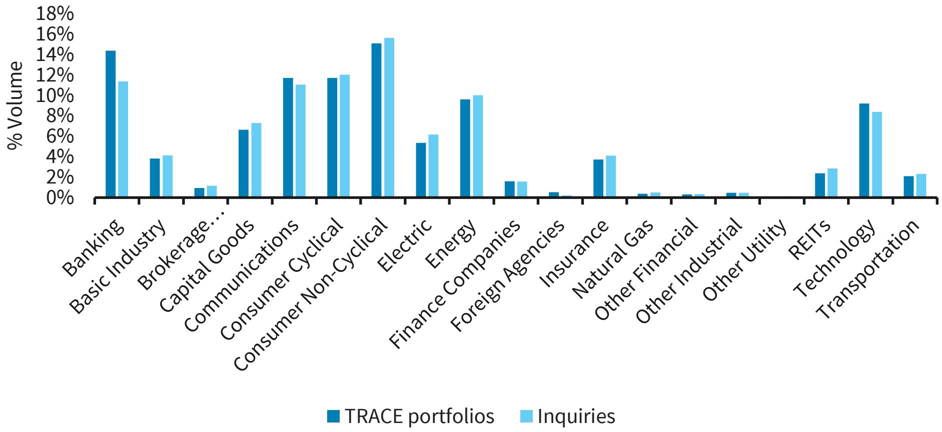

Figure 4: Distribution of Portfolio Volume by Liquidity Quintile

The figure shows the distribution of total IG portfolio trade volume by LCS quintile, where Q1 comprises the most liquid bonds and Q5 comprises the least liquid bond. Data based on portfolio trades identified by our ML algorithm executed during the period 1st January 2021 – 31st December 2021.

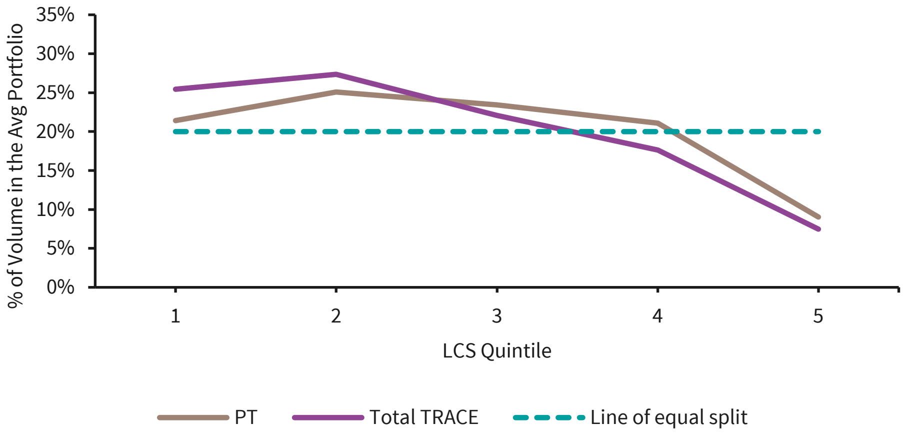

Figure 5: Mixing Liquid and Illiquid Bonds

The boxplot shows the distribution of the percentage of liquid volume (sum of trade volumes in the first two most liquid LCS quintiles) for Liquid and Illiquid portfolio trades. Liquid (Illiquid) PTs have lower (higher) trade volume-weighted LCS than the trade volume-weighted LCS of the bonds belonging to the Bloomberg IG Corporate Bond Index. Lower (higher) LCS is better (worse). Each box gives the 25th, median (red line) and 75th percentile of the distribution. Data based on portfolio trades identified by our ML algorithm executed during the period 1st January 2021 – 31st December 2021.

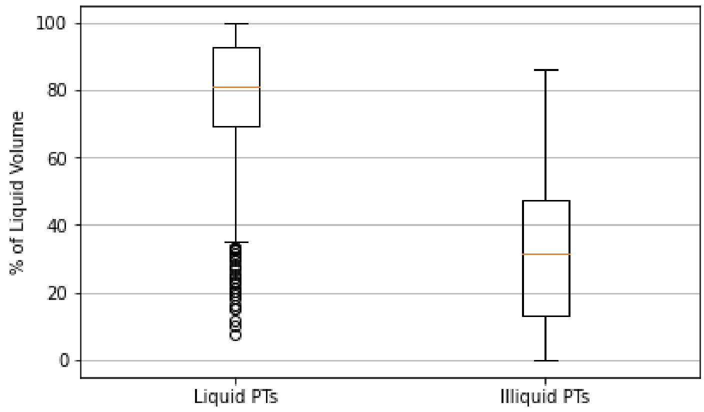

Figure 6: Overlap Between Portfolio Trades and ETFs

The figure shows the overlap between the line items of IG portfolio trades and the monthly holdings of LQD. An overlap of 0 means that *none* of the bonds in a given portfolio trade are held by LQD in that month; conversely, an overlap of 1 means that all of the bonds in the portfolio are held by LQD in that month. Data based on portfolio trades identified by our ML algorithm executed during the period 1st January 2021 – 31st December 2021.

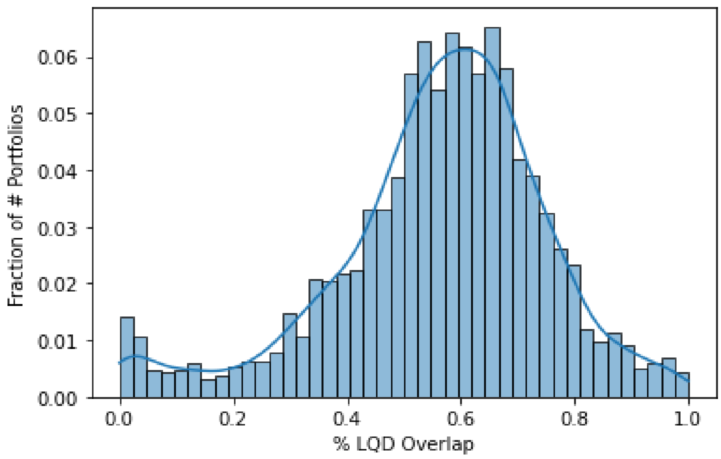

Figure 7: Transaction Cost Analysis – Identification Strategy

The figure shows the distribution of bond-date observations by trading protocol. For example, if a bond *j* traded only in RFQs on date *t,* this bond-date observation would be classified as "Only RFQs". If a bond *j* traded in both at least one RFQ and at least one PT, this bond-date observation would be classified as "Both PTs and RFQs". Data based on portfolio trades identified by our ML algorithm executed during the period 1st January 2021 – 31st December 2021.

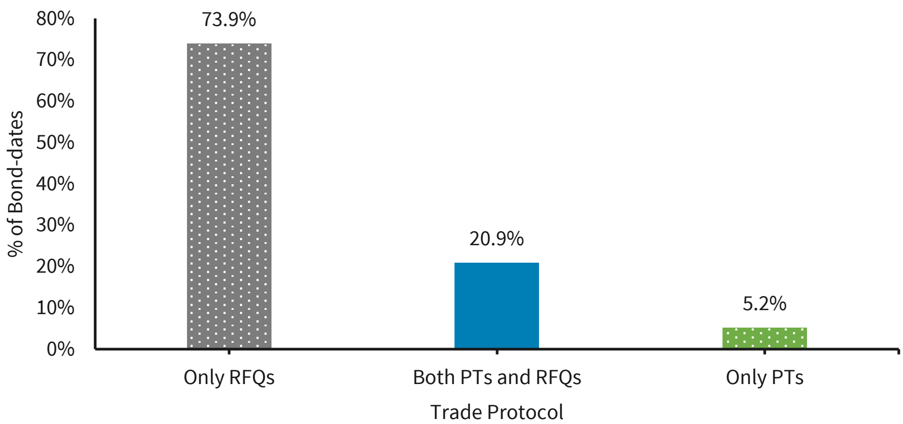

Figure 8: Portfolio Trade Execution Costs Over Time

The figure shows the difference between the average EHS of transactions executed in the PT and RFQ protocol (LHS) and the PT volumes as a percentage of TRACE (RHS). Smaller values of the difference indicate lower transaction costs of the PT protocol compared to the RFQ protocol. For more details on the definition of our TRACE universe refer to Section 2.3. Data are based on portfolio trades identified by our ML algorithm executed during the period 1st January 2020 – 31st December 2021.

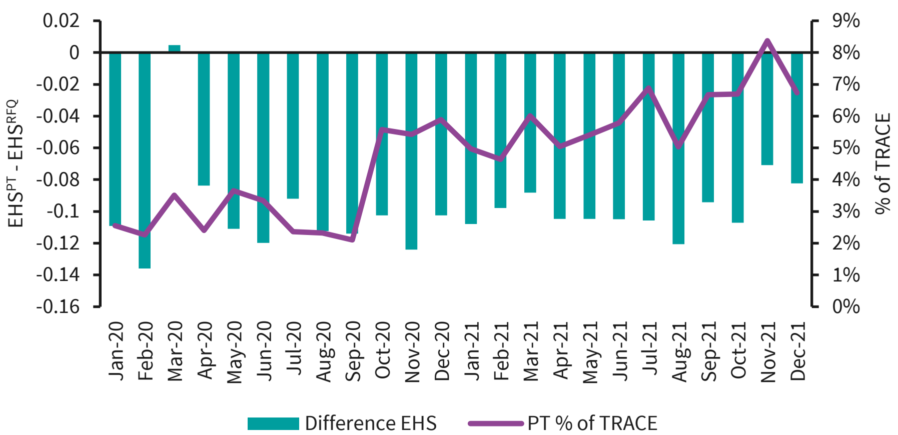

List of Tables

Table I: A Portfolio Trade Examples

The table gives an example of a portfolio inquiry that has 145 line items. The portfolio is identified by the *PT ID* and the line items that comprise that portfolio are identified by the *PT ID Line Item*. Note that since data are proprietary, all values displayed in the table are for illustrative purposes only and do not represent actual inquiries.

Date	PT ID	PT ID Line Item	CUSIP	Trade Size	Direction
2021-01-05	123	123_1	05971KAE9	$250,000	Buy
2021-01-05	123	123_2	03835VAG1	$500,000	Buy
2021-01-05	123	123_3	037833CJ7	$750,000	Buy
2021-01-05	…	…	172967LD1	…	…
2021-01-05	123	123_144	29444UBE5	$300,000	Buy
2021-01-05	123	124_145	404119BN8	$500,000	Buy

Table II: The Portfolio Trades Database

The table presents summary statistics of the portfolio trades database constructed using our ML algorithm. The estimate of the TRACE market excludes non-index corporate bonds, but includes volumes at common spotting times. For more details on the bond sample and discussion around spotting times, refer to Section 2 and Section 3.

	# Bond-PT Obs.	**# of PTs**	**$ Volume (bln)**	**% of TRACE**
		Panel A: Aggregate		
2018-2021	998,975	12,107	696	3.47
		Panel B: Time Series		
2018	107,541	1,950	81	1.14
2019	175,224	2,265	127	1.68
2020	245,774	2,978	177	3.09
2021	470,436	4,914	311	6.89

Table III: Algorithm Validation – True Positives and False Positives

The table shows how well the clustering algorithm is capable of identifying the "true" portfolios in the investor inquiries database. For any given inquiry, the true positives rate is calculated as the number of line items the algorithm identified divided by the total number of line items in that inquiry. The false positives rate is defined as the number of incorrectly identified line items divided by the total number of line items the algorithm found.

	True positive rate		False positive rate	
	In-sample 2021	Out-of-sample H1 2022	In-sample 2021	Out-of-sample H1 2022
Mean	85%	77%	15%	14%
Median	97%	100%	3%	7%

Table IV: Algorithm Validation – Empirical Distribution

The table compares the empirical distributions of the investor inquiries (INQ) and the portfolio trades (PT) identified using our ML algorithm along two dimensions: characteristics of the portfolio (Panel A) and characteristics of the bonds in the portfolio (Panel B). Portfolio LCS, Maturity and Age (time since issuance) are computed as a weighted average where the weights are given by the notionals of the line items in that portfolio. Our sample of inquiries comprises those inquiries which we could successfully match in full to the TRACE database. The last row, *TRACE ex PT,* reports the volume-weighted LCS, maturity and bond age for all non-portfolio trades in TRACE.

	Panel A: Portfolio Characteristics								**Panel B: Bond Characteristics**					
	# Line Items		Volume ($ mn)		Line Item Wgt (%)		# of Sectors		LCS (%)		Maturity (years)		Bond Age (years)	
	INQ	PT	INQ	PT	INQ	PT	INQ	PT	INQ	PT	INQ	PT	INQ	PT
Mean	93	97	76.3	68	2.16	2.04	11	12	0.83	0.84	9.44	10.22	2.53	2.62
Std	114.9	115.67	118.6	109	1.73	2.93	3.45	3.05	0.29	0.45	6.01	5.94	1.32	1.38
P25	27	37	14.1	21	0.83	0.97	9	10	0.68	0.59	6.15	6.21	1.66	1.71
Median	51	60	36.2	34.4	1.75	1.72	11	12	0.81	0.75	7.1	8.13	2.25	2.42
P75	109	105	89.9	69	3.12	2.78	14	14	0.92	0.97	10.66	12.72	3.12	3.17
TRACE ex PT	*NA*								0.69		10.7		2.44	

Table V: Share of Portfolio Trading by Illiquidity and ETF Ownership Quintiles

The table shows the average monthly share of portfolio trading as a percentage of total monthly trade volume for bonds double-sorted by LCS and ETF ownership. ETF ownership is computed using monthly portfolio holdings of all ETFs included in the CRSP Mutual Funds Database in 2021.

		ETF Ownership					
		Low	2	3	4	High	H-L
Illiquidity (LCS)	Low	5.9%	8.7%	8.3%	9.2%	9.7%	3.8%
	2	7.9%	6.8%	9.0%	10.5%	11.2%	3.3%
	3	5.7%	5.4%	7.1%	9.6%	7.9%	2.2%
	4	8.8%	9.5%	11.6%	12.1%	12.1%	3.3%
	High	11.2%	11.9%	12.8%	13.5%	14.9%	3.7%
	H-L	5.3%	3.2%	4.5%	4.3%	5.2%	-

Table VI: Transaction Costs of Portfolio Trades

$$\boldsymbol{EHS_{i,j,t} = \beta_1 Portfolio\ Trade_{i,j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t}}$$

The table reports transaction-level regressions of effective half spread ($EHS_{i,j,t}$) on a portfolio trade dummy ($Portfolio\ Trade_{i,j,t}$) and a set of controls. Regressions include the following transaction-level controls (collected in vector $Z_{i,j,t}$): a dummy equal to one if a trade is larger than \$5 million ($Block\ Trade_{i,j,t}$), previous trade *EHS* ($Lag\ EHS_{i,j,t}$), a dummy equal to one if the trade is executed before 13:00 EST, a dummy equal to one if all five previous trades were all investor sells ($Sell\ Pressure_{i,j,t}$) and a dummy equal to one if all five previous trades were investor buys. All continuous variables are winsorized at the 1% level. Regressions include either bond-date fixed effects ($\boldsymbol{\lambda_{j,t}}$) or more granular fixed effects based on bond, date, and size or direction. *Size FE* are based on \$0.5 million buckets (all trades above \$ 5 million belong to the same size bucket). *Direction FE* differentiate between dealer buys and sells. T-stats in parentheses. Standard errors are clustered at the bond and date level. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Panel A: Point Estimates

	TRACE Portfolios			Investor Inquiries		
	(1)	(2)	(3)	(4)	(5)	(6)
Portfolio Trade	-6.05*** (-23.47)	-6.00*** (-23.35)	-6.38*** (-32.67)	-5.24*** (-4.70)	-5.24*** (-4.70)	-6.14*** (-6.53)
Block Trade	-2.92*** (-11.26)	-	-	-2.62*** (-10.20)	-	-
Lag EHS	-0.16*** (-32.28)	-0.16*** (-32.26)	-0.08*** (-23.37)	-0.16*** (-32.31)	-0.16*** (-32.34)	-0.08*** (-23.43)
Morning Trade	-0.39*** (-3.32)	-0.40** (-3.39)	-0.59*** (-5.20)	-0.09 (-0.78)	-0.10 (-0.87)	-0.29*** (-2.60)
Sell Pressure	0.85** (2.52)	0.84** (2.53)	0.53*** (5.05)	0.88*** (2.65)	0.88*** (2.66)	0.53*** (5.05)
Buy Pressure	3.28*** (7.13)	3.30*** (7.17)	0.51*** (3.18)	3.36*** (7.30)	3.37*** (7.35)	0.53*** (3.36)
Bond-Date FE	YES	NO	NO	YES	NO	NO
Bond-Date-Size FE	NO	YES	NO	NO	YES	YES
Bond-Date-Size-Direction FE	NO	YES	YES	NO	YES	YES
Bond-trade Observations	4,853,114					
Sample Period	Jan 1st 2021 – December 31st 2021					

Panel B: Percentage Improvement

	(1)	(2)	(3)	(4)	(5)	(6)
Mean EHS RFQs	14.5bp	14.5bp	14.5bp	13.9bp	13.9bp	13.9bp
Improvement vs. RFQs	41.7%	41.4%	44.0%	37.7%	37.7%	44.2%

Table VII: Transaction Costs of Portfolio Trades By Illiquidity Profile

$$\boldsymbol{EHS_{i,j,t} = \beta_1 Port\,Trade_{i,j,t} + \beta_2 Port\,Trade_{i,j,t} \times Illiq\,Dummy_{j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t}}$$

The table reports transaction-level regressions of effective half spread on a portfolio trade dummy ($Port\ Trade_{i,j,t}$) and an interaction term with an illiquidity dummy ($Illiq\ Dummy_{j,t}$). $Illiq\ Dummy_{j,t}$ equals to one for bonds in the most illiquid quintile. Regressions include the following transaction-level controls (collected in vector $Z_{i,j,t}$): a dummy equal to one if a trade is larger than \$5 million ($Block\ Trade_{i,j,t}$), previous trade *EHS* ($Lag\ EHS_{i,j,t}$), a dummy equal to one if the trade is executed before 13:00 EST, a dummy equal to one if all five previous trades were all investor sells ($Sell\ Pressure_{i,j,t}$) and a dummy equal to one if all five previous trades were investor buys. All continuous variables are winsorized at the 1% level. All regressions include bond-date fixed effects ($\boldsymbol{\lambda_{j,t}}$). T-stats in parentheses. Standard errors are clustered at the bond and date level. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Panel A: Point Estimates					
	(1) LCS	(2) TES	(3) Bond Age	(4) Price Impact	(5) Roll
Portfolio Trade	-2.76 (-1.36)	-6.13*** (-20.80)	-3.54** (-1.95)	-3.53** (-2.08)	-5.29*** (-14.75)
Portfolio Trade ×Illiq Dummy	-5.21*** (-2.78)	-2.15*** (-6.38)	-3.23* (-1.89)	-7.09*** (-4.57)	-3.09* (-1.89)
Block Trade	-3.21*** (-10.58)	-2.96*** (-9.55)	-3.15*** (-10.23)	-3.15*** (-10.26)	-3.20*** (-10.64)
Lag EHS	-0.92*** (-9.65)	-0.24*** (-4.65)	-0.92*** (-9.65)	-0.92*** (-6.21)	-0.92*** (-9.65)
Morning Trade	-0.71 (-1.37)	-0.71*** (-4.85)	-0.69 (-1.35)	-0.69 (-1.34)	-0.70 (-1.35)
Sell Pressure	1.61*** (2.67)	0.84** (1.99)	1.62*** (2.68)	1.61*** (2.67)	1.62*** (2.68)
Buy Pressure	6.45*** (7.65)	4.20*** (6.61)	6.44*** (7.67)	6.45*** (7.68)	6.45*** (7.67)
Bond-Date FE	YES	YES	YES	YES	YES
Bond-trade Observations	4,853,114	4,394,447	4,853,114	4,853,114	4,853,114
Sample Period	Jan 1st 2021 – December 31st 2021				
Panel B: Percentage Improvement					
Mean EHS RFQs	14.5bp	14.5bp	14.5bp	14.5bp	14.5bp
Improvement vs. RFQs (Liquid)	19.3%	42.3%	24.4%	24.3%	36.5%
Improvement vs. RFQs (Illiquid)	55.2%	57.1%	46.7%	73.2%	57.8%

Table VIII: Portfolio-level Characteristics

The table reports summary statistics of $EHS_p - \widehat{EHS_p}$ (Panel A) and of the portfolio-level characteristics (Panels B and C) used in Section 6: *Corr LQD* (past 30-day correlation between portfolio and LQD returns), *Max Corr Top 10 ETFs* (the highest correlation with any of the top 10 largest IG ETFs), *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of IG ETFs), *% LQD C/R* (% of line items that were in either the create or redeem basket of LQD in the five days after the portfolio trade), *% Rightway LQD C/R* (% of line items that were "rightway" for the LQD C/R process on the day the portfolio was executed), *% Liquid Rightway LQD C/R* (% of total illiquid line items that were "rightway" for the LQD C/R process on the day the portfolio was executed), *% Illiquid Rightway LQD C/R* (% of total illiquid line items that were "rightway" for LQD C/R process on the day the portfolio was executed), and *Diversification* (one minus the ratio of the standard deviation of the portfolio returns and the weighted sum of the standard deviation of the returns of the individual line items).

	Mean	Std	P25	P50	P75
Panel A: Predicted Cost Benefit of Portfolio Trading					
$EHS_p - \widehat{EHS_p}$	-2.86bp	6.11bp	-6.69bp	-2.99bp	0.48bp
Panel B: Characteristics Linked to the ETF Ecosystem					
Corr LQD	0.85	0.09	0.83	0.92	0.94
Max Corr Top 10 ETFs	0.91	0.06	0.89	0.94	0.96
% ETF Ownership	6.4%	1.2%	5.6%	6.6%	7.4%
% LQD C/R	36.5%	17.2%	25.0%	35.9%	47.6%
% Rightway LQD C/R	8.4%	13.5%	0%	3.2%	11.2%
% Liquid Rightway LQD C/R	9.5%	15.7%	0%	2.8%	12.3%
% Illiquid Rightway LQD C/R	9.4%	15.8%	0%	1.9%	12.0%
Panel C: Characteristics Linked to Diversification					
Diversification	0.13	0.07	0.08	0.11	0.16

Table IX: Relationship to the ETF Ecosystem – Portfolio-level Analysis

$$\boldsymbol{EHS_p - \widehat{EHS}_p = \alpha + \beta_k Portfolio\ Characteristics_{k,p} + \epsilon_p}$$

The table reports the results of cross-sectional regressions of actual (*EHS*) minus predicted ($\widehat{EHS}$) effective half spread of portfolio $\boldsymbol{p}$ on portfolio-level characteristics: *Corr LQD* (past 30-day correlation between portfolio and LQD returns), *Max Corr Top 10 ETFs* (the highest correlation with any of the top 10 largest IG ETFs), *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of IG ETFs), *% LQD C/R* (% of line items that were in either the create or redeem basket of LQD in the five days after the portfolio trade), *% Rightway LQD C/R* (% of line items that were "rightway" for the LQD C/R process on the day the portfolio was executed), *% Liquid Rightway LQD C/R* (% of total illiquid line items that were "rightway" for the LQD C/R process on the day the portfolio was executed), and *% Illiquid Rightway LQD C/R* (% of total illiquid line items that were "rightway" for LQD C/R process on the day the portfolio was executed). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively. Panel A shows point estimates; Panel B shows the % change in $EHS_p - \widehat{EHS}_p$ (relative to its mean) given an interquartile shift in the dependent variables.

	$EHS_p - \widehat{EHS}_p$					
	Panel A: Point estimates					
	1st Link: ETF Pricing and Hedging			2nd Link: ETF C/R		
	(1)	(2)	(3)	(4)	(5)	(6)
Corr LQD	-3.20*** (-4.13)	-	-	-	-	-
Max Corr Top 10 ETFs	-	-2.72** (-1.97)	-	-	-	-
% ETF Ownership	-	-	-0.23** (-2.29)	-	-	-
% LQD C/R	-	-	-	-1.84*** (-3.13)		-
% Rightway LQD C/R	-	-	-	-	-2.37*** (-2.87)	-
% Liquid Rightway LQD C/R		-	-	-	-	-0.42 (-0.39)
% Illiquid Rightway LQD C/R		-	-	-		-1.93** (-2.14)
Sample	4,914 PTs	4,914 PTs	4,914 PTs	3,838 PTs	3,838 PTs	3,838 PTs
	Panel B: Effect of an interquartile shift					
Change in $EHS_p - \widehat{EHS}_p$, % (vs. mean)	12.2%	6.6%	14.5%	14.7%	9.4%	8.0%[21]

[21] The estimate is based on *% Illiquid Rightway LQD C/R*

Table X: Alternative Explanations – Diversification

$$\boldsymbol{EHS_p - \widehat{EHS}_p = \alpha + \beta_k Portfolio\ Characteristics_{k,p} + \epsilon_p}$$

The table reports the results of cross-sectional regressions of actual (*EHS*) minus predicted ($\widehat{EHS}$) effective half spread of portfolio $\boldsymbol{p}$ on portfolio-level characteristics: *Corr LQD* (past 30-day correlation between portfolio and LQD returns), *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of IG ETFs), *% Rightway LQD C/R* (% of line items that were "rightway" for the LQD C/R process on the day the portfolio was executed) and *Diversification* (one minus the ratio of the standard deviation of the portfolio returns and the weighted sum of the standard deviation of the returns of the individual line items). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively. Panel A shows point estimates; Panel B shows the % change in $EHS_p - \widehat{EHS}_p$ (relative to its mean) given an interquartile shift in the dependent variables.

	$EHS_p - \widehat{EHS}_p$	
	Panel A: Point Estimates	**Panel B: % change (vs. the mean) given an interquartile shift in the dependent variables**
Corr LQD	-4.95*** (-3.75)	19.0%
% ETF Ownership	-0.39*** (-3.91)	24.5%
% Rightway LQD C/R	-2.15*** (-2.64)	8.4%
Diversification	-4.25*** (-2.59)	11.8%
Sample	3,113 PTs	

Table XI: Alternative Explanations – Investors "Swapping" Portfolios

The table shows the probability of an offsetting portfolio trade (i.e. in the opposite direction as the original PT) covering at least 50% of the line items in the original trade happening on the same day (T) or in the next 5 business days (T+1, T+2, T+3, T+4 and T+5).

	T	**T+1**	**T+2**	**T+3**	**T+4**	**T+5**
Mean	0.0057	0.0037	0.0039	0.0041	0.0038	0.0042
Std	0.075	0.061	0.062	0.063	0.062	0.065
P25	0	0	0	0	0	0
Median	0	0	0	0	0	0
P75	0	0	0	0	0	0
Observations	4,914 PTs executed during the period Jan 1st 2021 – Dec 31st 2021					

References

Agarwal, V., Hanouna, P., Moussawi, R., & Stahel, C. (2018). Do ETFs Increase the Commonality in Liquidity of Underlying Stocks. *SSRN Abstract 3001524*.

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

Ben-David, I., Franzoni, F., & Moussawi, R. (2018). Do ETFs increase volatility? *Journal of Finance, 73*(6), 2471-2535.

Bessembinder, H. (2003). Issues in assessing trade execution costs. *Journal of Financial Markets, 6*, 233-257.

Bessembinder, H., Jacobsen, S., Maxwell, W., & Venkataraman, K. (2018). Capital commitment and illiquidity in corporate bonds. *Journal of Finance, 73*(4), 1615-1661.

Carapella, F., & Monnet, C. (2020). Dealers' insurance, market structure and liquidity. *Journal of Financial Economics, 138*(3), 725-753.

Choi, J., Huh, Y., & Seunghun Shin, S. (2023). Customer liquidity provision: Implications for corporate bond transaction costs. *Management Science*.

Choi, J., Kronlund, M., & Oh, J. (2022). Sitting Bucks: Zero Returns and Stale Pricing in Fixed Income Funds. *Journal of Financial Economics, 145*(2), 296-317.

Collin-Dufresne, P., Junge, B., & Trolle, B. (2020). Market Structure and Transacation Costs of Index CDSs. *The Journal of Finance, 75*(5), 2719-2763.

Corwin, W., & Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance, 67*, 719-760.

Da, Z., & Shive, S. (2018). Exchange traded funds and asset return correlations. *European Financial Management, 24*, 136-168.

Dannhauser, C., & Hoseinzade, S. (2022). The unintented consequences of corporate bond ETFs: Evidence from the taper tantrum. *The Review of Financial Studies, 35*(1), 51-90.

Dick-Nielsen, J. (2009). Liquidity Biases in TRACE. *The Journal of Fixed Income, 19*(2), 43-55.

Dick-Nielsen, J. (2014). How to Clean Enhanced TRACE Data. *Technical Report*. Retrieved from http://ssrn.com/abstract=2337908

Dick-Nielsen, J., Feldhütter, P., & Lando, D. (2012). Corporate Bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics, 103*(3), 471-492.

Edwards, A., Harris, L., & Piwowar, M. (2007). Corporate Bond Market Transaction Costs and Transparency. *The Journal of Finance, 62*(3), 1421-1451.

Ester, M., Kriegel, H.-P., & Sander, J. X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In E. Simoudis, J. Han, & U. Fayyard (Ed.), *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining.*

Feldhütter, P. (2012). The same bond at different prices: identifying search frictions and selling pressures. *Review of Financial Studies, 25*, 1155-1206.

Friewald, N., Jankowitsch, R., & Subrahmanyam, R. (2012). Illiquidity or credit deterioration: a study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics, 105*(1), 18-36.

Goldberg, J., & Nozawa, Y. (2020). Liquidity Supply in the Corporate Bond Market. *Journal of Finance*.

Goldstein, I., Jiang, H., & Ng, T. D. (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics, 126*, 592-613.

Goldstein, M., & Hotchkiss, E. (2020). Providing liquidity in an illiquid market: Dealer behavior in US corporate bonds. *Journal of Financial Economics, 135*(1), 16-40.

Hagströmer, B. (2021). Bias in the effective bid-ask spread. *Journal of Financial Economics, 142*(1), 314-337.

Holden, C., & Nam, J. (2019). Market accessibility, Corporate bond ETFs, and Liquidity. *Available at SSRN 3083257*.

Jiang, H., Li, D., & Wang, A. (2021). Dynamic Liquidity Management by Corporate Bond Mutual Funds. *Journal of Financial and Quantitative Analysis, 56*(5), 1622-1652.

Konstantinovsky, V., Yuen Ng, K., & Phelps, B. (2016). Measuring Bond-Level Liquidity. *Journal of Portfolio Management, 42*(4), 116-128.

Koont, N., Ma, Y., Pastor, L., & Zeng, Y. (2022). Steering a Ship in Illiquid Waters: Active Management of Passive Funds. *SSRN Working Paper Abstract 4053844*. Retrieved from Available at SSRN: https://ssrn.com/abstract=4053844

Lesmond, D., Ogden, J., & Trzcinka, C. (1999). A new estimate of transaction costs. *Review of Financial Studies, 12*, 1113-1141.

Li, J., O'Hara, M., Rapp, A., & Zhou, A. (2023). Bond Market Illiquidity: Is Portfolio Trading the Solution? Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4495516

Marta, T. (2020). Fixed Income ETFs, Bond Liquidity, and Stressed Markets. *Available at SSRN 3350519*.

Meli, J., & Todorova, Z. (2022, July 22). Portfolio Trading: Part III- In Crisis and Calm. *Barclays Research Paper*.

Meli, J., & Todorova, Z. (2023). Credit ETFs in Mutual Funds and Corporate Bond Liquidity. *Financial Markets, Institutions & Instruments, 32*(3), 89-114.

O'Hara, M., & Zhou, A. (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics, 140*(2), 368-390.

Pu, X. (2009). Liquidity commonality across the bond and CDS Markets. *Journal of Fixed Income, 19*, 26-39.

Rhodes, M., & Mason, J. (2022). ETF ownership and firm-specific information in corporate bond returns. *Journal of Financial Markets*.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance, 39*, 1127-1139.

Schestag, R., Schuster, P., & Uhrig-Homburg, M. (2016). Measuring liquidity in bond markets. *Review of Financial Studies, 29*, 1170-1219.

Shim, J., & Todorov, K. (2021). ETFs, illiquid assets and fire sales. *SSRN Working Paper Abstract 3886881*. Retrieved from https://ssrn.com/abstract=3886881

Ye, S. (2019). How do ETFs Affect the Liquidity of the Underlying Corporate Bonds? *Chinese University of Hong Kong Working Paper*.

Data Appendix
A1. The Machine Learning Algorithm
Step 1: Matching Barclays Inquiries to TRACE

The TRACE rules require dealers to report a trade for each individual bond in the portfolio with an attributed dollar price and an execution timestamp, despite the fact that technically the dealer and the client agree on a single price for the entire basket of bonds. This means that the individual line items must appear in TRACE if a portfolio inquiry is executed. For each line item in our portfolio inquiries database, we search through the TRACE database for Dealer-to-Customer trades which exactly match the line items in the inquiry on CUSIP, date, quantity traded and direction (dealer buy or dealer sell) [22]. The result of this process is a database of *traded* inquiries, augmented with the *exact* execution time stamp and the executed price, both of which are recorded in TRACE.

In 85% of the matches we find in TRACE, there is exactly one trade which satisfies the criteria above. The difficulty comes from the remaining 15% for which there are multiple matches. This occurs because our inquiries database only records the date but not the exact execution time stamp. Due to the enormous number of trades in TRACE, in some cases it is not possible to identify the trade without the time stamp. This applies particularly for trade sizes less than $250K and trades executed around busy times sometimes cannot be identified without the exact execution timestamp. To determine the most likely candidate for a given bond belonging to a portfolio inquiry where multiple candidates are available, we use the distribution of the execution timestamps of the other line items in that portfolio trade. For example, we know that on Jan 5th 2021, a $250K buy trade in bond X was part of a portfolio inquiry. Assume we find three such trades executed at 10:00, 12:30 and 14:30. If the majority

---

[22] We define a dealer as a traditional broker-dealer or as an alternative trading system (ATS), which we identify setting the field "Reporting Party Type" as either "D" or "T". A dealer must report a trade if the counterparty is either a customer or an affiliate, which we identify by setting the field "Contraparty Type" as either "C" or "A".

of the other line items in that inquiry were executed around 14:30, we would select the $250K trade in bond X at 14:30 and discard the other two candidate trades.

To create the blueprint of the typical portfolio trade, we need to convert the matched line items to matched portfolio inquiries. However, in doing so, we face the following trade-off. On the one hand, we want to find as many of the inquiries that actually traded as possible, but on the other, we want to minimize the number of individual bond trades we incorrectly classify as part of a portfolio i.e. the false positives. To strike the optimal balance between these goals we can pull two levers – (1) the maximum time period within which those line items must be executed; and (2) the minimum percentage of line items in the inquiry required to classify a portfolio as found[23].

To illustrate, assume that we require to find at least 80% of the line items in an inquiry. As we increase the time span between the trades that we consider, we will match more of the line items, and thus match more of the portfolios. However, we also risk over-classifying trades in TRACE, which just happen to have the same notional and the same direction as the portfolio inquiry but were not part of it. ***Figure A1.1*** demonstrates this trade-off. If we allow a time interval of 2 hours, we find 70% of the portfolio inquiries, but we over-classify 20% of the line items (i.e. 20% of the line items have multiple matches). By tightening the time interval to 5 minutes, we steeply reduce over-classification to c.2% at the cost of finding only slightly fewer of the inquiries. We conclude that 5 minutes is the optimal time interval since tightening beyond that only marginally improves precision, but drastically reduces the proportion of the inquiries that we can find.

[23] For example, if we were only able to match two line items out of one hundred included in a portfolio inquiry, and they occurred hours apart, then we clearly have not actually found the portfolio inquiry in TRACE. In contrast, if we find all one hundred line items within two seconds, then we are quite confident that we've located the portfolio trade.

Figure A1.1: Varying the Maximum Time Interval When Matching Portfolio Inquiries

The figure shows the trade-off between the percentage of portfolio inquiries we find in TRACE against the over-classification error, as we vary the maximum time span we allow between the first and the last line item in any given portfolio inquiry.

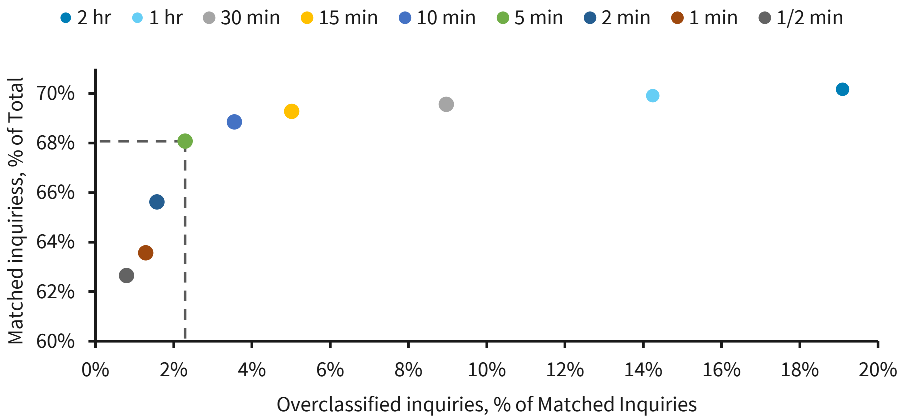

In ***Figure A1.2*** we vary the threshold of line items per portfolio we require to match in TRACE. We find 68% of all inquiries with at least 80% of line items. In comparison, we find 57% of all inquiries with at least 95% of line items. Interestingly, the percentage of inquiries we find decreases from 100% to 75% as we just increase the threshold from 0% to 5%, but then decreases only very slowly as we further tighten the criterion. This suggests that we either find the inquiries (almost) in full or we don't find them at all. This confirms the anecdotal evidence we have received from our conversations with portfolio desk traders about the take-it-or-leave-it nature of portfolio trades. Nonetheless, we set a rather conservative threshold of 80% of line items found in order to minimize classification error.

Figure A1.2: Varying the Minimum Number of Matched Line Items

The figure shows the percentage of portfolio inquiries we identify in TRACE as we vary the minimum number of matched line items per portfolio inquiry.

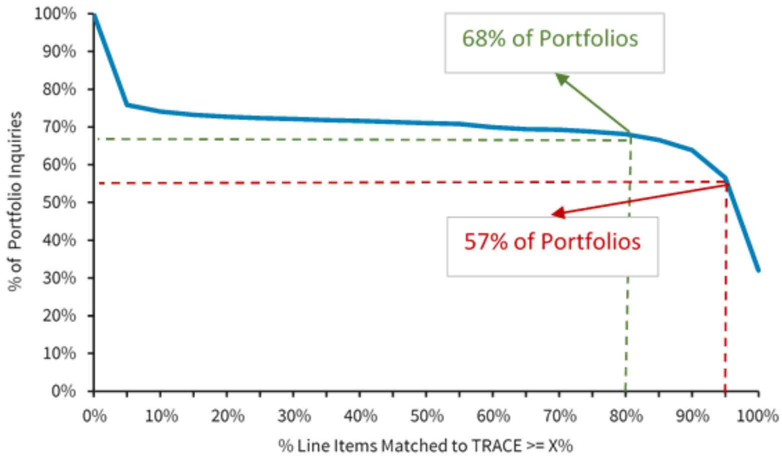

Step 2: Clustering algorithm

The two most important parameters of the ML algorithm we train are the maximum time we allow to elapse between the lines items in any given portfolio trade and the characteristics of the typical portfolio trade. We select and tune both parameters based on the proprietary dataset of portfolio inquiries matched to TRACE in Step 1.

Analysing the inquiries, we discovered that when we plot the number of trades recorded in TRACE per each second of the trading day, seconds during which portfolio inquiries were executed appear like spikes or clusters (***Figure A1.3***). However, the problem is that such clusters are both rare (compared to the total volume that appears on TRACE) and could take very different time to appear on the TRACE tape. For example, most inquiries span zero to two seconds, but some of the larger ones could take up to 20 seconds. This means that we need to develop an algorithm which is able to separate the large amount of "noise" in the data (i.e. the non-portfolio trades), but is flexible enough to accommodate different portfolio structures. In other words, the algorithm needs to allow for different length (in terms of time) of the portfolios. For example, an algorithm which identifies clusters based on a fixed time interval, no matter how tight that interval is, would produce noisier estimates.

We employ a machine learning toolkit and use a DBSCAN clustering algorithm (Density-based Spatial Clustering of Applications with Noise) to obtain a list of portfolio *candidates* (Ester, Kriegel, & Sander, 1996). DBSCAN searches through the millions of TRACE observations and forms clusters of trades whose execution timestamps are closely packed together (i.e. the trades have many nearby neighbours) and marks as outliers points that are located in low-density regions (i.e. their nearest neighbours are too far away). Clusters identified in this way are strictly non-overlapping and the individual line items included in any cluster are unique to that cluster only.

More specifically, each day from January 1st 2018 to December 31st 2021, the algorithm orders all dealer-to-customer trades in TRACE by their timestamp and, starting from the first trade on that day, searches for trades which have at least 25 other trades recorded within a two second interval. Each such a trade is labelled as a "core" trade. Some of the trades that are within the two second neighbourhood of a "core" trade could be "core" trades themselves. For example, if a trade that is exactly two seconds from the original "core" trade has at least 25 trades within its own two second window, it too would be a "core" trade. We then link "core" trades and their two second neighbourhoods to form a cluster. In other words, each cluster must contain at least one core point. Further, individual trades in a cluster may well be more than two seconds apart, but *any* trade in a cluster is *at most* two seconds away from some core trade. It is precisely the "expanding" nature of the algorithm, which produces clusters with different time length.

Figure A1.3: Examples of How Portfolio Trades Appear in TRACE

Example 1 “Tight” – Inquiry with 244 Bonds

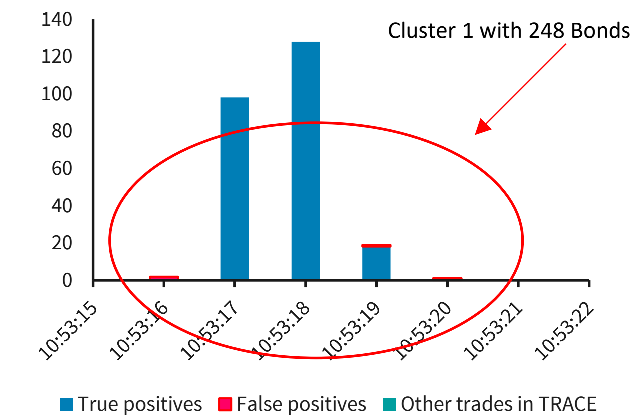

Example 2 “Spread-out” – Inquiry with 105 Bonds

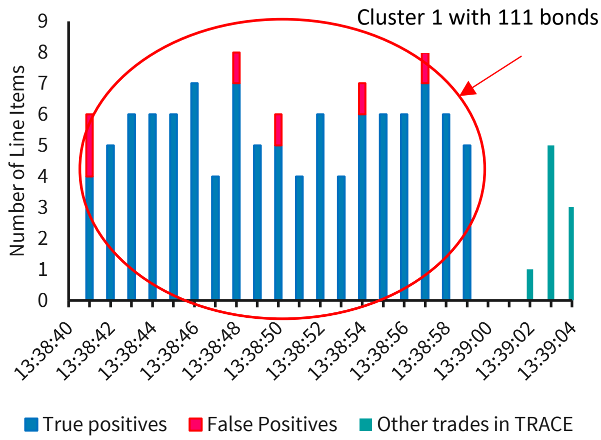

Example 3 “Batched” – Inquiry with 98 Bonds

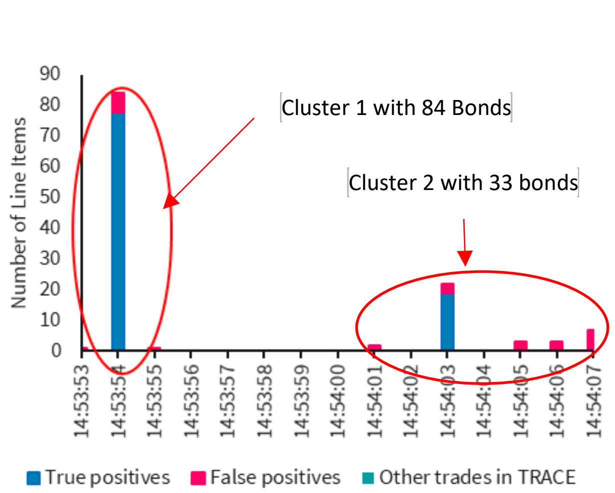

Example 4 “Batched” – Inquiry with 88 Bonds

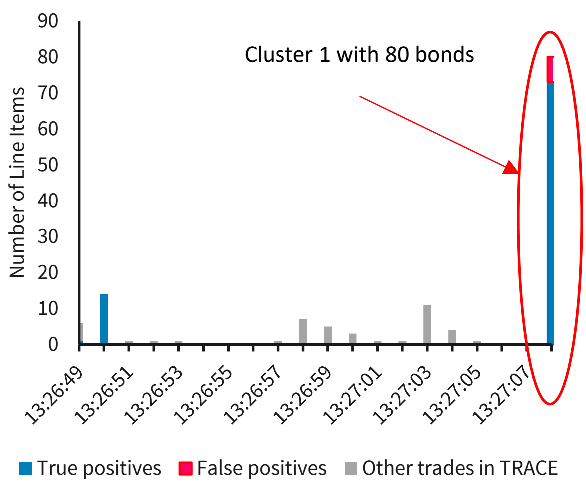

Step 3: Re-clustering

Each of the clusters the algorithm identifies has a high probability of being an actual portfolio trade. As shown by the examples above, some portfolio trades are split into multiple batches, and others are in one, and it is extremely difficult for our algorithm to tell which is which. This is why we think this approach will give accurate estimates of the volumes associated with this trend, but less accurate estimates of the boundaries of portfolio trades, and hence the overall count. However, since we eventually want to test how execution quality differs across different portfolio construction strategies, it would be extremely valuable to reconstitute these clusters into their original portfolios, if possible.

To this end, we re-classify the clusters from the previous step by aggregating those clusters that happen one minute apart into a unified portfolio. The idea is to bring together several batches of the same portfolio (as in Example 3 and 4 on ***Figure A1.3***). It is important to note that we neither add nor delete portfolio volume in this step – we simply adjust the boundaries of the clusters.

Step 4: Filtering

Next, we filter this list to remove candidate clusters than don't line up with what we expect given the analysis of our inquiry in Step 1:

- We drop clusters that are within 5-minute intervals before and after popular delayed spot times – 11.00, 15.00, 15.30, 16.00, 16.30. As a result, we are likely to understate the true prevalence of portfolio trades because some IG portfolio trades are certainly spotted at these times. However, if we don't drop those clusters we are certain to include lots of false positives.
- We drop clusters that contain less than \$5 million in HY and \$10 million in IG, and clusters with average line item size below \$100K in HY and below \$250K in IG. This is necessary to reduce false positives associated with odd lots trading, much of which is electronic.

Finally, we fine-tune by deleting a small number of line items which are markedly different from the cluster to which they belong. These adjustments have a minor impact on the total portfolio volume we identify but substantially reduce the portfolio-level false positives rate:

- We know from our inquiries that portfolio trades are either buy-only, sell-only or balanced buy-and-sells trades. For example, if we see a candidate cluster with 100 bonds, the most likely distribution of trades is – a client buys 100 bonds; a client sells 100 bonds and client buys 50 bonds and sells 50 bonds. Hence, a candidate cluster where a client buys 97 bonds and sells 3 bonds is extremely unlikely – in reality this is a buy-only trade with 97 bonds and the 3 sell trades were coincidentally executed at the same time.

- Similarly, the majority of portfolio inquiries are HY only or IG only. Whenever HY and IG bonds are mixed in the same portfolio trade, these are likely to be at the boundary between HY and IG – e.g. a mix of BAA3s and BA1s. In other words, a candidate cluster of 95 B2 bonds and 5 AAA bonds is highly unlikely, even if the direction (buy or sell) matches. This is likely to be a straight HY trade with 95 bonds.

A2. Figures

Figure A2.1: Correlation Between Actual and Implied ETF Flows

The figure shows the histogram of monthly correlation coefficients between actual and implied LQD flows. We obtain actual flows data from Bloomberg. Implied flows data are estimated following the procedure developed by Shim and Todorov (2021) and Koont et al. (2022). Estimation period January 2018-December 2022.

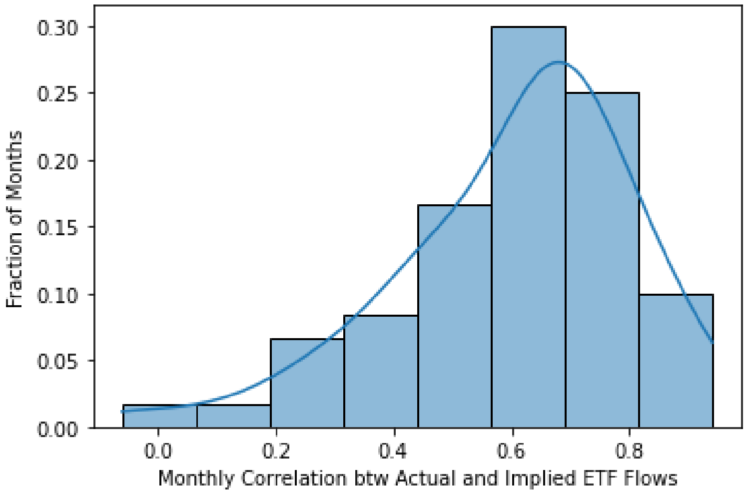

A3. Tables

Table A3.1: Other Portfolio Strategies

The table shows a summary of the different strategies investors use when trading portfolios of bonds. When classifying a portfolio trade we always compare it to the Bloomberg IG Corporate Bond Index. A *Liquidity* strategy applies in those cases where the portfolio trade contains at least 50% more trade volume in the 4th or 5th LCS quintile (most illiquid quintiles) than what we would normally expect in the IG Index. A *Market View* strategy applies in those cases when the portfolio maturity/sector/rating Herfindahl score (HHI) is at least 50% higher than the respective HHI of the Index. Portfolio trades that are neither axed towards a *Liquidity* nor a *Market View* strategy are classified as *Diversified*.

Type of Strategy	**% of IG PT Volume**
➢ Liquidity	49%
➢ Market View	
o Maturity View	35%
o Sector View	24%
o Rating View	13%
Concentrated (Liquidity OR Market View)	**69%**
Diversified (Flows Management)	**31%**