Portfolio Trading in Corporate Bond Markets

Jeffrey Meli (a)

NYU

Zornitsa Todorova (b)

Barclays, UK

27 June 2025

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 construct a database of portfolio trades from TRACE and demonstrate that they have 40% lower transaction costs than trades in individual bonds. Trade-level and portfolio-level analysis reveal that two linkages between ETFs and portfolio trades drive the reduction in transaction costs. ETFs enable market-makers to efficiently price and hedge portfolio trades and provide an additional outlet for the risk accumulated via portfolio trading.

JEL Codes: C55, G12, G14

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

We thank Itzhak Ben-David, Rene Stulz, Adam Kelleher, Arik Ben Dor, Carlo Favero, Claudio Tebaldi, Daniele d'Arienzo, Melissa Prado, Nicholas Hirschey, an anonymous referee, and seminar participants at the 2023 AFA, Columbia Business School, Ohio State University, Bocconi University and Nova School of Business and Economics for helpful comments.

(a) Corresponding author: Jeffrey Meli, Tel: +1 646 765 5621, email: *jm11241@stern.nyu.edu*, Kaufman Management Center, 44 W4th Street, NY, NY 10012

(b) Zornitsa Todorova: 1 Churchill Place, London, ENG E14 5HP, zornitsa.todorova@barlcays.com

1. Introduction

Portfolio trading is a new protocol in the corporate bond market, in which an investor bundles a set of individual corporate bonds into one basket and executes it as a single piece of risk, with one market-maker. In this article we use a novel dataset of portfolio trades ("PTs") that we construct from TRACE to demonstrate that PTs are remarkably effective; they reduce realized transaction costs by more than 40% compared to single security trades ("SSTs"), with the bulk of the gains accruing to less liquid bonds. We link these reduced transaction costs to the corporate bond ETF ecosystem. ETFs provide market-makers with real-time price transparency and hedging capability for ETF-like portfolios of corporate bond risk and an additional outlet for the risk market-makers accumulate through portfolio trades. 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. We expose a new channel through which corporate bond ETFs affect their underlying securities: ETFs help overcome the constraint on corporate bond liquidity imposed by the vast number of unique CUSIPs. Unlike equities, a bond issuer can have many securities outstanding, each of which has a different maturity and coupon, and possibly different seniority, optionality, and covenants. This 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. Liquidity provision by market-makers is limited because they may fail to find the other side of the trade, which results in holding a difficult-to-hedge position in inventory.

PTs turn this constraint on bond liquidity into a strength. When executing a PT, market-makers no longer need to find an exact and immediate offsetting trade for each CUSIP. Instead, they can price and hedge the transaction using ETFs and work out of the risk over time, by trading with other investors and using the ETF creation and redemption process. Put differently, PTs are an elegant solution to the matching problem that weighs on bond liquidity. Market-makers pass on the associated benefits to investors in the form of lower transaction costs, as they do when executing "agency", or "matchmaking" trades, which have low transaction costs because dealers take no inventory risk (Goldstein & Hotchkiss, 2020). The protocol is useful to investors because they naturally own many (often thousands) of CUSIPs, and there are many aspects of corporate bond portfolio management that can be accomplished, or even require, trading large numbers of bonds. In addition to index-like exposures, PTs are used to adjust curve, sector, or credit quality exposures. Investors can now execute these types of trades via PTs at significantly lower transaction costs than when executing them as a series of individual trades.

We demonstrate the effectiveness of portfolio trading using a rigorous regression specification, where we compare the difference in the realized transaction cost incurred when trading the same bond on the same day in a PT versus an SST. We control for trade-level characteristics and include bond-date fixed effects. We find that PTs have 40% lower transaction costs than SSTs. The main concern with this result is related to causality, which arises if the choice of protocol is endogenous. Our result is biased if specific types of transactions are both more cost effective and more likely to be executed via PTs. We address this by isolating trades associated with index rebalancing, defined as investor buy (sell) trades in bonds that are added (dropped) from the most common investment grade ("IG") and high yield ("HY") bond indices, in a tight window around month-end. These transactions are driven by exogenous demand to trade, and share a common motivation, urgency, and direction, and thus the remaining variation in trading protocol is exogenous. The results are identical to those based on the full sample.[1]

The applicability and attractiveness of the PT protocol is evidenced by its rapid adoption; PTs grew from 1% of total corporate bond volume at their inception in 2018 to 7% in 2021. While their popularity with investors is understandable, it remains unanswered why market-makers are willing to execute PTs at such low transaction costs. The high overlap between the bonds owned by the largest ETFs and those included in PTs strongly suggests a role for ETFs. For example, 56% of the line items in IG PTs are owned by LQD, one of the largest IG ETFs, whereas LQD owns only 30% of the bonds in the Bloomberg IG Corporate Bond Index. Conversations with practitioners reveal that many market-makers have re-organized the layout of their trading floors to co-locate the PT and ETF trading operations, which speaks to the synergies between the two products. Finally, we demonstrate that the reduction in transaction costs associated with PTs is not uniformly distributed across bonds. The greatest benefit accrues to less liquid bonds and to bonds that are heavily owned by ETFs.

We identify two ways market-makers can leverage the ETF ecosystem to efficiently execute PTs. First, market-makers can take advantage of the real-time pricing and deep secondary market liquidity of bond ETFs (Meli and Todorova, (2023)) to price and hedge

[1] See Dick-Nielsen and Rossi (2019) and Ottonello (2019), amongst others.

portfolios. While there may be significant uncertainty about the price of any one individual bond, ETFs sharpy reduce the uncertainty about the value of an ETF-like portfolio of bonds. Second, ETFs provide market-makers an additional outlet for the risk they accumulate through PTs, both directly, though the creation and redemption process, and indirectly, through increased investor transactions in bonds owned by ETFs.

Careful portfolio construction is essential to activate these channels; one of our key results is that PTs that are "closer" to ETFs get better execution. To demonstrate the importance of these channels, and the associated constraints on portfolio construction, we create a portfolio-level measure of the benefit of transacting via a PT. We use our baseline regression model to predict the transaction cost of each line item in a PT if it had traded as an SST and define the "PT benefit" as the difference between the actual PT transaction cost and this predicted SST cost. We then take the value-weighted average of this measure across the line-items to arrive at the portfolio-level PT benefit.

Equipped with this measure, we show that the benefit of executing via a PT is larger for portfolios that are easier to price and hedge using ETFs, proxied by the correlation of recent returns and the overlap with the largest ETFs. For example, an interquartile shift in the ETF ownership of bonds (as a percentage of the outstanding amount) results in 16% better execution for HY PTs. High ETF ownership is also associated with greater volumes and trade counts, even after controlling for bond liquidity, suggesting that market-makers can more easily close out positions in those securities. In IG, PTs over-index for bonds that will be included in future C/R baskets, even after adjusting for the fact that they over-index for bonds owned by those ETFs, and the PT benefit is larger for portfolios with a greater overlap with ETF C/R baskets. An interquartile shift in the percentage of a portfolio that is right way vis-à-vis the C/R process increases the benefit of executing an IG PT by 16%.

Activation of these channels imposes significant constraints on portfolio construction. At a basic level, ETFs are only useful for pricing and hedging portfolios; given their diversified nature, ETFs are of limited use for pricing and hedging transactions in individual bonds. In addition, market-makers must be able to map the real-time secondary market prices of ETFs into the price of the portfolio. This requires that portfolios be sufficiently diversified across sectors and issuers. It also requires that portfolios contain a sufficient proportion of liquid bonds, with recent-enough transactions to facilitate this mapping. Market-makers would struggle to map the price of an ETF into the price of a portfolio of bonds that never trade.

Activation of the second channel requires that the illiquid portion of the portfolio be heavily owned by ETFs, such that market-makers are confident that they can offload positions in those bonds. The ability to include some illiquid bonds is an important feature of portfolio trades. The pricing and hedging benefits conferred by the more liquid bonds effectively allow investors to “crowd-source” liquidity in securities that are otherwise difficult and expensive to trade. However, the benefits only accrue to the subset of illiquid bonds that are heavily owned by ETFs.

These requirements are also the main constraints on the use of PTs. The protocol is not well-suited for trades in individual securities, issuers, or narrow market segments, because ETFs are not useful for pricing and hedging such portfolios. PTs are also most effective for portfolios comprised primarily of bonds that are heavily owned by ETFs, explaining why they over-index for these bonds. We expect that SSTs will remain the dominant trading protocol in those cases for which PTs are not effective, which (at least currently) make up most trading needs. In addition, the protocol is unlikely to work well when ETF flows are one sided. For example, there was no PT benefit in the IG market during the covid-induced volatility of March 2020. IG bond ETFs experienced sustained redemptions (in early March) followed by sustained creates (in late March). Most PTs were in the same direction as the ETF flows, which deactivated the create and redeem channel. Given these constraints, we see the PT protocol as an efficient alternative available that investors will utilize in certain circumstances, whist relying on trades in individual securities for their other trading needs.

We consider and reject alternative explanations for the reduced transaction costs associated with PTs, including that PTs limit adverse selection for market-makers, that the gains are spurious due to misallocation of bond prices, and that PTs are effective because they are diversified and thus impose less idiosyncratic risk on market-makers. The fact that PTs over-index for bonds that are not heavily owned by ETFs is itself evidence against the adverse selection and misallocation hypotheses; given the informational spillovers generated by ETFs, these bonds are the least likely to benefit from reduced adverse selection, and to have stale mid prices that market-makers can utilize to create a spurious PT benefit.

Our analysis of portfolio trading is based on a PT database that we construct, which is necessary because PTs were not flagged in the TRACE feed until May 2023.[2] We build this

[2] The reporting rule changed on the 15th May 2023, as described in FINRA’s *Regulatory Notice 22-12*. The rapid growth of the PT protocol was the main motivation for the addition of the flag.

database using a proprietary set of PT inquiries received by a large market-maker. We match these inquiries to TRACE to find a verified set of executed PTs. 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. The resulting dataset contains more than 16,000 unique PTs and c.1.4 million bond-PT transactions. We perform a number of validation checks to ensure that the algorithm identifies actual portfolio trades. This algorithm is a contribution in its own right; it facilitates study of PTs from their inception until May 2023, and can be applied to markets (e.g., the European corporate bond market) where trade reporting does not include a PT flag.

Relationship to prior literature

Our analysis contributes to several areas of the existing literature. First, we contribute to the literature on the evolution of the supply of and demand for corporate bond liquidity 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 agency 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 focused on investors' response 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 most recent stage in the development of new trading protocols and the management of liquidity needs. PTs significantly reduce transaction costs, and their wide applicability to credit investors, who naturally own and trade large numbers of CUSIPs, is behind the rapid uptake of the new protocol.

We also contribute to the literature on the implications of ETFs. In equities, the literature shows that ETFs have a positive effect on volatility (Ben-David et al. (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 ETF-like portfolios of corporate bond risk. 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, even though the investors utilizing the protocol need not ever buy or sell an ETF directly. This also speaks to the limitations of the protocol, which is limited to trading strategies that can be implemented trading large ETF-like portfolios.

Our article is also related to work by Li et al. (2023), who apply our algorithm to study PTs. The authors argue that PTs are effective because they diversify market-maker inventories and thus reduce idiosyncratic risk. We distinguish between the diversification hypothesis and our ETF hypothesis using long-short PTs: those in which the investor both buys and sells bonds. These are an ideal sub-sample because the ETF linkages still apply. We expect long-short PTs will be cost effective so long as each leg can be individually priced and hedged using ETFs. In contrast, a long-short PT contains mostly (or entirely) idiosyncratic risk, and thus they increase, rather than decrease, the idiosyncratic exposure of the market maker. We find that long-short PTs have the same costs as long-only or short-only PTs. Further, those long-short PTs that are close to balanced, meaning that the long and short legs are of roughly similar notional have identical execution as those that are unbalanced, in direct contrast to the prediction of the diversification hypothesis.

2. Data and Variables Definitions

2.1 Portfolio Trades vs. Single Security Trades

In a portfolio trade an investor requests a quote on a portfolio of corporate bonds as a single piece of risk. If the investor agrees to the price, the portfolio trade is executed in its entirety with a single market-maker. This bundling of transactions in individual bonds into one large trade distinguishes a PT from the standard SST protocol, through which an investor executes trades in individual bonds. Investors at times execute many SSTs at once, where the quotes are obtained via “Wanted In Competition” lists (known as “WICs”) of bonds to multiple market-makers.[3] The individual trades that comprise these lists are typically referred to as line items 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, as they are in a PT.

In all other respects, execution under the PT and SST protocols is the same. In both, the investor submits requests for quote (“RFQs”) to market-makers, the standard procedure through which investors obtain executable prices for a specific trade. In both protocols investors balance the potential for information leakage against the desire to obtain the best execution when choosing which and how many market-makers to send their RFQs. Finally, both types of transactions are reported to TRACE (Trade Reporting and Compliance Engine).

However, there is one nuance associated with TRACE reporting of PTs: although a single price is agreed to for the entire portfolio, trade reporting is done at the individual line-item level. We assume that the line-item prices reported to TRACE are accurate (i.e., that the portfolio price is correctly apportioned across line items), allowing us to compare the line-item level transaction costs of PTs to those of SSTs. This assumption is potentially problematic. While the prices of the individual line items reported to TRACE (weighted by their respective notional) must sum to the quoted price of the portfolio, a market-maker seemingly has discretion as to how it allocates that price across line items. For example, the market-maker may inflate the prices of bonds it was long and deflate the prices of those it was short, or misprice a bond in order to generate investor interest in it.

There are several significant constraints on market-makers’ practical ability to engage in this type of strategic misallocation. First, convention requires that market-makers provide

[3] These are typically “BWICs” (bids wanted in competition) or “OWICs” (offers wanted in competition).

both the aggregate portfolio price and the prices of the individual line items when responding to a PT RFQ. Investors must know both before agreeing to a trade because they have best execution requirements that apply at the bond, as opposed to the portfolio, level. An investor who bought a portfolio where some trades were priced too richly would attract scrutiny, regardless of whether other bonds in the same portfolio were priced cheaply. Similarly, market-makers are subject to fair dealing requirements. A market-maker that reported a bond purchase to TRACE at an artificially low price could be accused of excessive mark-up if it then sold the same bond at the correct price to a different investor.

These constraints reduce the scope for misallocation; any mispricing must be too small to be noticed or of concern to the investor or regulators. Second, even before the addition of the PT flag in TRACE in 2023, market participants could infer that a trade was part of a PT by observing the clustering of trades in TRACE; if market-makers routinely mispriced bonds in PTs investors would ignore that print, rendering the mispricing useless. Finally, any small mispricing that does occur despite these constraints will add noise, not bias, and will not affect our results given our sample size.

There is one important caveat that we must consider. Most PTs contain some bonds that only trade in PT form; these tend to be illiquid and thus may have stale mid prices. Identification in our baseline empirical specification is driven by bonds that trade in both protocols on a given day. It is possible that market-makers could take advantage of stale marks on illiquid line items to create a spurious reduction in transaction costs, whereby bonds with easily observable mid prices have attractive quotes, but those with stale mid prices appear to be quoted fairly but are in fact expensive vis-à-vis true mid. We explicitly rule out this possibility in Section 6, where we consider alternative explanations for our findings.

2.2 Portfolio Trade Inquiry

Although a transaction level comparison of PTs and SSTs is conceptually possible, PTs were not flagged in TRACE until May 2023, over five years after they began trading. This presents a practical challenge to any such analysis. To create a dataset of PTs, our first step is to collect a proprietary dataset of investment grade (IG) and high yield (HY) 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 Barclays executed the trade. The inquiry dataset contains c.3,000 investor inquiries and c.22,000 unique corporate 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, and direction (buy or sell). **Table *I*** gives an example of a typical inquiry.[4]

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 volume in 2021 (**Figure *1***). While we believe that Barclays' market share is sufficient to ensure 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 comprehensive database of portfolio trades. Hence, in ***Section 3***, we use the proprietary inquiries database to develop an algorithm to identify portfolio trades which are not included in our inquiry data.

2.3 Bond Sample and Liquidity Measures

We obtain transaction-level corporate bond data from the standard version of TRACE, which caps trade sizes at $5 million for IG bonds and at $1 million for HY bonds, for the period 1st October 2018- 31st December 2021.[5] We follow Dick-Nielsen ((2009), (2014)) to remove double counting, corrections, reversals and cancellations from TRACE. We 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.

[4] The dataset also contains a flag if the portfolio is "custom" or "in-competition". We discard the custom portfolios (10% of the 2021 sample and 20% of the 2018-2020 sample) because they are designed by the market maker 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.

[5] 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.

Our main measure of liquidity is Liquidity Cost Score (LCS), 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). For robustness purposes, we compute four additional liquidity metrics at the bond-month level: bond age, Trade Efficiency Score (TES), Price Impact, and Roll's measure. Including both quotes-based and trade-based measures allows for robust measurement of liquidity for bonds that trade infrequently, and abstracts from differences in realized execution costs linked to the use of "agency" (or "matchmaking") bond trading, which has increased (Kargar et al. (2021)). The details of these measures are included in the Appendix.

2.3 ETF Ownership and Creation and Redemption (C/R) Baskets

We obtain monthly ETF holdings from CRSP Mutual Funds Database, which we aggregate across ETFs to compute the ETF ownership as a percentage of amount outstanding for each CUSIP. In some tests, we also use daily holdings, which we then use to construct the daily ETF C/R baskets of the iShares iBoxx Investment Grade ETF (ticker: LQD) and the iShares iBoxx High Yield ETF (ticker: HYG), large IG and HY ETFs, respectively. These daily ETF holdings are publicly available and can be obtained from the iShares website.[6] Following the methodology by Shim and Todorov (2021) and Koont et al. (2022), we impute the realized creation and redemption baskets from daily changes in holdings on days with C/R activity. We identify create (redeem) days as those days on which there was a positive (negative) change in the number of LQD/HYG shares. We then use daily changes in the number of bonds held to infer the composition of the net LQD or HYG basket on each day.[7]

3. Constructing a Database of Portfolio Trades

We use the subset of our database of PT inquiries that traded (i.e., those we can find in the TRACE feed) to develop a "blueprint" of PTs in TRACE.[8] PTs appear as clusters of

---

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

[7] 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. Our imputed baskets are best interpreted as the average net basket on each given day. We have verified that the average (monthly) correlation between the actual and imputed ETF flows is close to 0.80.

[8] The untraded PT inquiries could in principle be informative about optimal portfolio construction. However, only a subset of these were legitimate inquiry; others were tests submitted by investors learning about the protocol. Unfortunately, these are not consistently labelled and thus we discard the non-executed trades.

trades with the same or very similar execution time stamps. We use this insight to tune a machine learning clustering algorithm to identify candidate PTs in the TRACE data. Another insight from the inquiry data is that the line items included in PTs share certain characteristics, including similar size, credit rating, and tenor. Accordingly, we develop a set of filters designed to reduce the rate of false positives included in each PT (given the tight window in which the line items of a PT are reported to TRACE, false negatives are rare). These steps are summarized in ***Figure 2***.

The resulting database contains over 16,000 unique portfolio trades, totalling c. 1.4 million bond-portfolio observations and $937 billion of executed bond volumes. We perform an extensive series of validation checks to confirm that we have identified actual PTs, including in-sample and out-of-sample tests for accuracy and precision, times-series and cross-sectional comparisons of the database we construct to the traded inquiries, and confirmation that our baseline empirical results are unchanged when we limit our sample to the traded inquiries. See the Appendix for a detailed description of the clustering algorithm, the filters we utilize, and the validation tests we perform.

3.1 PT Summary Statistics

We segment the PT database into IG and HY PTs, based on the credit rating of the line items in each portfolio (**Table *II***). We use a 50% cutoff (i.e., we label a PT "IG" if more than 50% of the line items are investment grade), but very few portfolios are truly mixed. For example, on average 92% of the line items in IG portfolios are rated IG. Our database contains 12,107 unique IG PTs and 4,157 unique HY PTs.

In ***Table III***, 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). The average portfolio trade contains c.100 line items and $70 million worth of notional, approximately equally split between line items.

We compare the PTs identified by our algorithm to SSTs on volume-weighted liquidity (LCS), maturity, and bond age (for SSTs see the last row *"TRACE ex PT"* of **Table III**). Along maturity and age, PTs are very similar to SSTs. In contrast, PTs are substantially less liquid than SSTs. The average LCS of the line items in PTs is 0.84% compared to 0.69% for SSTs (higher LCS implies lower liquidity). Further, this is not driven by a few very illiquid portfolio trades. More than 50% of the portfolios, both by count and by volume, have lower liquidity than the average SST. However, PTs are not exclusively comprised of illiquid bonds. Instead, the reduced average liquidity reflects a shift in the portfolio distribution towards less liquid bonds (notably towards the third and fourth quintiles of the liquidity distribution), but with substantial representation of all levels of liquidity (***Figure 3***).

A second distinguishing feature of PTs is that they over-index for bonds in the largest and most liquid ETFs (***Figure 4***). In Panel C and Panel D, we look at the overlap with LQD and HYG. On average, 56% of the bonds in IG portfolio trades are owned by LQD and 72% of the bonds in HY portfolios are owned by HYG. By comparison, LQD owns c.30% of the bonds in the BBG IG Index and HYG owns c.60% of the bonds in the Bloomberg HY Corporate Bond Index (BBG HY Index). The overlap grows when we consider a larger set of liquid ETFs. In Panel B we show the overlap with the largest ETFs (tickers LQD, VCIT, VCSH, IGIB, IGSB in IG and HYG, JNK, SHYG and SJNK in HY). On average 87% of the line items included in PTs are owned by the largest ETFs. This is notable because trading volumes in the corporate bond ETF market are highly concentrated in these ETFs.

These two stylized facts stand somewhat in contrast. ETFs tend to own liquid bonds, yet PTs over-index for both ETF bonds and for illiquid bonds. This suggests that PTs are particularly concentrated in the least liquid bonds owned by ETFs. To verify, for each bond-month in our sample, we perform a double-sort by liquidity (based on LCS) and the percentage of the amount outstanding of each bond that is owned by ETFs in that month and compute the proportion of total volumes that are executed using PTs in each bucket (using ownership across the full universe of bond ETFs in the CRSP database). The share of volumes executed using PTs increases as liquidity declines and as ETF ownership increases (Panel A (IG) and Panel B (HY) ***Table IV***). For example, for very liquid IG bonds with low ETF ownership, only about 6% of their trading volume is done via portfolio trades. For the least liquid IG bonds with high ETF ownership, PTs make up 15% of volume.

4. Bond-level Analysis: PTs Transaction Costs, Liquidity, and ETFs

We first analyze PTs through the lens of the individual line-items. Following the literature (e.g., Bessembinder (2003); Collin-Dufresne, Junge, & Trolle (2020); Hagströmer

(2021))[9], we measure the transaction cost 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 from the mid-price; smaller *EHS* implies lower transaction costs realized by investors.[10] We then multiply *EHS* by 10,000 to compute transaction costs in basis points (bp).

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 is not representative of current PTs. Therefore, we restrict our analysis to data from 2021, to study portfolio trading in a more mature stage of its development (the one exception being our analysis of PTs during March 2020).

4.1 Baseline Analysis: PTs Reduce Transaction Costs by Over 40%

We compare the transaction costs for PTs and SSTs in a formal regression model at the transaction level:

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

where $PT_{i,j,t}$ is a dummy variable equal to one when transaction *i* in bond *j* on date *t* is part of a portfolio trade (based on the PTs identified by our clustering algorithm). The main coefficient of interest in ***Model 1*** is β_1, which is the difference between the *EHS* of PTs and SSTs. If portfolio trading is cost-effective, we expect $\beta_1 < 0$. Our baseline specification includes bond-date fixed effects ($\lambda_{j,t}$). Therefore, identification comes from variation in the transaction costs of those bonds which trade in both protocols on a given day. We see 21% of the bond-date

[9] Using equities quotes data, Hagströmer (2021) shows that the EHS measured relative to the mid-price could overstate the true spread. While Hagströmer's result may apply to corporate bonds, it does not compromise our results since our interest is the difference in execution costs between PTs and SSTs, rather than the level.

[10] 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.

observations in our sample in both protocols, which is a meaningful portion and supports the empirical validity of our results (***Figure 5***).

The use of bond-date fixed effects implies that we need only include transaction level controls, 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 ($Block_{i,j,t}$).[11] We also include the lagged *EHS* ($Lag\ EHS_{i,j,t}$) to account for any momentum effects in bond-level transaction costs. 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 the dummy variables $Sell\ Pressure_{i,j,t}$ and $Buy\ Pressure_{i,j,t}$ equal to one if trade *i* and the five trades preceding trade *i* are all sells or buys respectively. We cluster standard errors at the bond and date levels to account for correlation over time within a given bond and across bonds on a given date.

We estimate the model separately for IG and HY bonds (***Table V***). PTs are substantially more cost-effective than SSTs ($\beta_1 < 0$) (columns (1) and (3)). All else equal, the average transaction cost of a line item in an IG portfolio trade is 6.1bp cheaper than the same trade executed in SST form. In HY, portfolio trades are 11.2bp cheaper than SSTs. In percentage terms, the transaction costs of IG (HY) PTs are 42.1% (41.7%) lower than SSTs.[12]

4.2 Causality

Due to our use of bond-date fixed effects, our main causality concern is at the transaction level: that certain types of transactions are both more likely to be executed as PTs and naturally have lower transaction costs. If this were true, then the PT designation could be a sorting mechanism (between trades with higher and lower transaction costs) rather than an

[11] 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. To ensure that we are not conflating a size effect with a PT effect, we also run a version of the baseline on trades that are less than \$5 million, using a continuous size control, and get nearly identical results.

[12] In columns (2) and (4) of **Table *V***, we re-estimate ***Model 1*** but define $PT_{i,j,t}$ using the sample of PT inquiries. The magnitude of the coefficients, their statistical significance, and the percentage improvement in transaction costs over the SST protocol remain unchanged. This is both a robustness check on the baseline result and one validation of the algorithm we use to identify PTs. The Appendix contains further robustness checks with bond-date-size and bond-date-size-direction fixed effects, to compare 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 and significance of the PT dummy is unchanged.

efficient trading protocol. The causality concern is particularly acute because the PT protocol is best suited to trading strategies that naturally involve many securities; these may have different liquidity characteristics than issuer- or security-specific trades.

To address this concern, we need to isolate trades that share a common motivation, time scale, and sense of urgency, such that the remaining variation in the choice of protocol is exogenous. We follow Dick-Nielsen and Rossi (2019), Ottonello (2019) and Dick-Nielsen, Poulsen and Rehman (2023) and consider inclusions and exclusions from the Bloomberg US Corporate Bond Index (IG Index) or the Bloomberg US High Yield Corporate Bond Index (HY Index) as a source of exogenous trading demand. Rebalancing is also a good example of the type of transaction level issue that could bias our results. Given the predictability of index rules, it is possible that market-makers enter month-end prepared to facilitate rebalancing trades, and therefore they are executed at low transaction costs. Rebalancing is also a natural fit for PTs: it involves trading large numbers of line items at small size. Therefore, these trades may be both efficient and likely to be PTs, making this an effective test for causality.

Both the IG and HY indices include dollar-denominated corporate bonds with more than one year to maturity, in addition to several other requirements, such as minimum par amount outstanding and seniority, and both are rebalanced on the last trading day of each month.[13] The rules for bonds entering and exiting the index are fully programmatic, transparent, and available to all market participants. The bonds that enter the IG Index are either newly issued or they are upgraded from high yield to investment grade; bonds that enter the HY Index are either newly issued or are downgraded from investment grade to high yield. Bonds exit the indices due to the aforementioned ratings changes, if the remaining time to maturity drops below the minimum threshold of one year, or if they are called by the issuer.

Index rebalancing generates exogenous demand to trade from passive bond funds, which track their respective benchmark indices: they buy bonds which enter the index and sell bonds which exit the index.[14] Since passive funds minimise their index tracking error, they have a strong motive to trade close to the rebalancing date. We isolate index rebalancing trades by limiting our sample to trades where investors buy (sell) bonds which enter (exit) the

[13] For details on the BBG IG Index eligibility criteria see *https://US-Corporate-Index.pdf* and *https://US-Corporate-High-Yield-Index.pdf* for the BBG HY Index.
[14] Meli and Todorova (2023) show that the majority of IG bond funds are indexed to the Bloomberg US Aggregate Index (BBG AGG Index), which includes IG bonds, Treasuries, and agency and mortgage-backed securities. Importantly for our analysis, the IG Index is fully contained in the BBG AGG Index.

IG Index or the HY Index during the three days[15] before and after the index rebalancing date. These trades share the same motivation, are done in the same direction, and share a common time frame for execution; thus, the choice of protocol can be disentangled from any liquidity considerations investors might have during their normal trading activity. We re-estimate ***Model 1*** using this subsample of trades, using data on both TRACE portfolios and investor inquiries, and report the results in **Table *VI***. In both IG and HY, we obtain very similar results to the baseline in **Table *V***. Executing IG (HY) rebalancing trades in PT form reduces transaction costs vis-à-vis SSTs by 42.5% (42.2%), compared to a PT benefit of 42.1% (41.7%) using the full sample of trades. Given the similarity of the results in both IG and HY, we are confident in the validity of the original baseline analysis.

4.3 PTs Are More Effective for Illiquid Bonds and ETF Bonds

Informed by our finding that PTs over-index for illiquid bonds and bonds owned by ETFs, we augment ***Model 1*** with two interaction terms to examine how the benefit of PTs varies across bonds based on those characteristics:

$$EHS_{i,j,t} = \beta_1 PT_{i,j,t} + \beta_2 PT_{i,j,t} \times Illiq\ Dummy_{j,t} + \beta_3 PT_{i,j,t} \times ETF\ Own_{j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t} \quad (\textbf{\textit{Model 2}})$$

The dummy variable $Illiq\ Dummy_{j,t}$ is equal to one for bonds which belong to the most illiquid quintile of LCS distribution (see the Appendix for specifications using other measures of liquidity). Therefore, $\beta_2 < 0$ indicates that portfolio trading is more effective for illiquid bonds. Similarly, the dummy variable $ETF\ Own_{j,t}$ is equal to one for bonds in the lowest quartile (half) of ETF ownership in the IG (HY) universe.[16] The coefficient β_3 captures any difference between the PT benefit of those "non-ETF" bonds and the rest of the universe; $\beta_3 > 0$ indicates that portfolio trading is less effective for bonds with limited ETF ownership.

Table *VII* (columns (1) and (3)) contains the results of ***Model 2***. In both IG and HY, the coefficient β_1 is negative, significant, and close in magnitude to the coefficient in the baseline specification. For example, in IG β_1 is 5.5bp, compared to 6.1bp in the baseline specification. In addition, illiquid bonds in both markets benefit to an even greater extent

[15] The results are unchanged using a narrower time frame around index rebalancing dates.
[16] We do not include $Illiq_{j,t}$ or $ETF\ Own_{j,t}$ as standalone variables in ***Model 2*** as their effect is subsumed by the bond-date fixed effects**.** We defined the IG and HY ETF dummies differently to reduce the collinearity with the illiquidity measure, as generally bonds that are less heavily owned by ETFs are also less liquid.

from PT execution ($\beta_2 < 0$). In IG, a bond in the lowest LCS quintile has an additional 3.1bp (= β_2) reduction in EHS associated with PT execution; in HY, the additional reduction in EHS is 6.2bp. Each roughly implies that illiquid bonds have a 60% increase in the benefit associated with PT execution. In the appendix we demonstrate that this result holds across each of our alternative measures of liquidity.

In addition, β_3 is positive and significant in both markets. Non-ETF bonds benefit less from PT execution. For a bond that is not classified as illiquid, being in the bottom quartile (half) of ETF ownership reduces the benefit of PT execution by 44% (12%) in IG (HY).[17] Together, these results both explain our earlier finding that PTs over-index for illiquid bonds and bonds owned by ETFs (these are precisely the bonds that benefit the most from PT execution) and provide the first evidence linking the effectiveness of PTs to ETFs.

Finally, we construct a measure of the portfolio-level weight of non-ETF bonds. For each line-item in a given PT, the variable $Port\ ETF_{i,j,t}$ is equal to the notional-weighted proportion of the portfolio that is in the bottom quartile (half) of ETF ownership in IG (HY) and is equal to zero for SSTs. We include both this variable itself and its interaction with $ETF\ Own_{j,t}$:

$$EHS_{i,j,t} = \beta_1 PT_{i,j,t} + \beta_2 PT_{i,j,t} \times Illiq\ Dummy_{j,t} + \beta_3 PT_{i,j,t} \times ETF\ Own_{j,t} + \beta_4 Port\ ETF_{i,j,t} + \beta_5 Port\ ETF_{i,j,t} \times ETF\ Own_{j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t} \quad (\textbf{\textit{Model 3}})$$

The coefficient β_5 describes how the PT benefit of a trade in a non-ETF bond is affected by the presence of other non-ETF bonds.[18] The results are contained in **Table *VII*** (columns (2) and (4)); β_5 is positive and significant in both IG and HY. In other words, the PT benefit of a trade in a given non-ETF bond declines as more non-ETF bonds are included in the portfolio.

We conclude that portfolio construction is an important determinant of the effectiveness of the PT protocol. With a well-constructed portfolio, investors can achieve efficient execution on some non-ETF bonds. For example, based on the results in **Table *VII***, a singular non-ETF bond in an IG portfolio retains 72% of the baseline PT benefit, and an even greater proportion in a HY portfolio. The linkages to the ETF ecosystem are retained so long as these bonds are not too prominent in the portfolio because their effect is diluted by the large number of line items. However, as the proportion of non-ETF bonds grows, they

---

[17] As a robustness check, we double sort trades by liquidity and ETF ownership quintiles; in line with the above, the lowest quintiles of ETF ownership had the lowest average PT benefit.
[18] The construction of *Port ETF* obviates the need to interact it with the PT dummy.

begin to disrupt the linkages to ETFs, and their collective execution becomes less efficient. A one standard deviation increase in the proportion of non-ETF bonds reduces the PT benefit of each of those bonds by 12% in both IG and HY.

5. Portfolio-level Analysis: The Relationship to the ETF Ecosystem

We identify two linkages to the ETF ecosystem that drive the reduction in PT execution costs. First, ETFs provide market-makers an intra-day hedging and pricing tool for transactions in ETF-like portfolios of corporate credit risk. Market-makers can price a portfolio by examining its characteristics relative to liquid bond ETFs and use their real-time prices to adjust the price of the portfolio vis-a-vis the most recent marks, traded prices, and/or mid-prices of the individual bonds. Then, rather than search for an offsetting trade in each line-item in the portfolio, the market-maker can hedge its portfolio by taking offsetting positions in the optimal mix of ETFs, locking in the real-time ETF prices. In contrast, the pricing of a single security will reflect mostly idiosyncratic risks and security-specific supply and demand. Second, ETFs generate both direct and indirect transactions in the securities they own, through the C/R process and increased investor flows. These provide market-makers an alternative outlet for risk they accumulate through PTs, which is particularly important for illiquid bonds, where potential for the position to sit in inventory is high.

The uncertainty about current market prices and the difficulty hedging and offloading single-security bond positions are two costs that are reflected in the traditionally high bid-offer of principal trades in corporate bonds. That these costs matter is clear from the rise of "agency" (or "matchmaking") trading, in which market-makers line up both sides of the trade before executing. These trades have become more frequent since the financial crisis and have significantly lower bid-offer than principal trades precisely because they circumvent the costs associated with pricing and hedging (Goldstein & Hotchkiss, 2020). Like with agency trades, we posit that the reduction in these costs is reflected in the lower bid-offer spreads of PTs.

Motivated in part by the importance of portfolio construction at the line-item level, we assess the importance of these ETF channels by identifying the drivers of the PT benefit at the portfolio level. To construct a portfolio-level estimate of the PT benefit we need the counterfactual: where the portfolio would have traded as a series of SSTs. The first input in this estimate is the predicted half-spread, $\widehat{EHS_{\iota,J,t}}$, 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 $\widehat{EHS_{\iota,J,t}}$ 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 (value-weighting is necessary because the PT benefit is not uniformly distributed across bonds). 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, via a series of cross-sectional regressions:

$$EHS_p - \widehat{EHS_p} = \alpha + \beta_k Portfolio\ Characteristic_{k,p} + \epsilon_p \text{ (}\textbf{\textit{Model 4}}\text{)}$$

We control for portfolio-level liquidity (using the notional-weighted LCS) in ***Model 4***, as the line-item analysis suggests that less liquid portfolios will have a larger PT benefit.

5.1 IG PTs: Hedging and Pricing and the C/R Process

In IG, we measure the importance of the pricing transparency and hedging flexibility afforded by ETFs using the correlation of the portfolio with ETFs. We create a variable *Corr ETF*, 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. For one-directional portfolios (i.e., long-only or short-only), we compute the (absolute) correlation directly at the (global) portfolio-level; for long-short portfolios, we compute the (absolute) correlation with each leg separately and then take the notional-weighted average to arrive at a portfolio-level measure. We take this step because the ability to price and hedge each leg, rather than the aggregate portfolio, is the relevant measure (we rely on this distinction below when we consider alternative explanations for the PT benefit). With an average correlation of 0.79, LQD is an effective price benchmark and hedge for many IG portfolios (***Table VIII***).

Portfolios which are more effectively priced and hedged using ETFs incur lower transaction costs (*Panel ETF Pricing and Hedging*, ***Table IX***). Increasing the correlation of an IG PT from 71% to 93% (an inter-quartile shift) results in a 6% improvement in execution. Market-makers incur less basis risk when transacting portfolios that are more correlated to ETFs and pass (at least some of) the associated savings onto investors.

Market-makers can also use the ETF C/R process to offset positions acquired via PTs, by delivering the bonds purchased to create ETF shares (or the reverse for bonds sold). Bonds that are heavily owned by ETFs have a higher probability of being included in the daily ETF create or redeem baskets than bonds with low ETF ownership.

We define *% ETF C/R* as the percentage of the line items in each portfolio that are in the imputed creation or redemption basket of LQD during the five business days after the portfolio was executed. On average, 35.9% of IG PT line items are in the weekly C/R baskets of LQD (***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. For example, we confirm that the weekly LQD C/R baskets contain on average 40% of the total ETF holdings. Given that on average 56% of the bonds in a PT overlap with the holdings of LQD, we would expect an average overlap of 22.4% (= 56%*40%) between the portfolios and the C/R baskets, substantially below the actual level. In other words, PTs over-index for bonds 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 on the specific day of the PT significantly increases the improvement in transaction costs from executing via a PT (***Table IX***). We define a trade in an individual bond that an investor sells (buys) to a market-maker as "right way" for LQD when the bond is part of the create (redeem) basket that day, and aggregate this at the portfolio level in the variable *% Rightway ETF C/R*.[19] An interquartile shift in the percentage overlap with the C/R baskets of LQD (i.e., a shift from 0% to 13.8%, ***Table VIII***) increases the benefit of executing via a PT by 16%. We also create a version of this variable that tracks the percentage of the portfolio that was included in C/R baskets over the next week (*% ETF C/R*); an interquartile shift increases the PT benefit by 7%. This is based on the future (realized) C/R baskets, implying 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).

5.2 COVID and the ETF C/R Channel

As a further test of the ETF C/R channel, we examine the performance of IG PTs during the initial bout of COVID-induced volatility. This is an interesting period to study because the flows in IG corporate bonds were one-sided; investors were initially heavy sellers of IG

[19] This limits our sample for two reasons. First, the methodology by Shim and Todorov (2021) cannot define the ETF C/R baskets for each day, and second, imputed C/R flows tend to be economically very small (close to 0%) for many days in our sample. This is why we limit our sample to days with significant ETF C/R activity, defined as absolute imputed fund flows of LQD larger than 0.5%.

corporate bonds and then, once the Federal Reserve announced its emergency facilities, heavy buyers (in contrast, HY flows were more balanced in this period). In keeping with these aggregate flows, the large IG ETFs experienced a sustained period of redemptions and then a sustained period of creations. For example, the shares outstanding in LQD declined by nearly 2% between March 10th and March 19th and then increased by over 30% from then until March 31st. In both periods the flows were consistently one-sided.

Most PTs during this period were in the same direction as the aggregate flows; investors tried to transact in any way they could. Therefore, the ETF C/R channel was inactive. Market-makers could not offload the risk they accumulated via PTs to ETFs because they were being asked to buy (sell) portfolios of bonds while the ETFs were redeeming (creating) shares. The absence of the C/R channel should reduce or eliminate the PT benefit. To test this, we run our baseline model on IG PTs that were executed between March 10, 2020 and March 31, 2020 (**Table *X***, column (1)). The coefficient on the PT dummy is positive and insignificant. As expected, there was no reduction in transaction cost associated with PTs at that time (and possibly a slight penalty).

This result reflects the fact that most PTs were in the same direction as the aggregate flows. However, there were some “contrarian” PTs where investors were buying bonds during the sell-off or selling them during the rally. These could take advantage of the C/R channel, and thus it is possible that they did achieve superior execution to SSTs. To test this, we create a dummy variable *Against Market*, which is equal to 1 for trades that were “right way” vis-à-vis the aggregate flows (e.g., the investor was buying the bond in early March). We include both that variable and its interaction with the PT dummy in an augmented version of ***Model 1*** (**Table *X***, column (2)). The coefficient on the PT dummy in this specification is positive and significant; PTs that were aligned with the aggregate flows got worse execution than similar SSTs. However, the coefficient on the interaction term is negative, significant, and larger in magnitude than the coefficient on the PT dummy. PTs that went against the prevailing flows (for which we need to sum the two coefficients) not only got better execution than PTs that were aligned with the prevailing flows but achieved better execution than similar SSTs (the coefficient on the standalone variable *Against Market* is insignificant). This provides further evidence of the importance of the ETF C/R channel. Where this channel is inactive, PTs are not efficient vis-à-vis SSTs and may even be inefficient.

5.3 HY PTs: Hedging and Pricing and ETF Ownership

IG bonds trade frequently (relative to HY bonds), even those that are illiquid. Therefore, we believe our estimated correlation is an accurate measure of the true correlation between IG portfolios and ETFs. Further, most IG bonds have limited idiosyncratic risk, and co-movement of returns is the major risk that market-makers incur when executing a PT. Neither of these is true in HY, however. Many illiquid HY bonds do not trade frequently, and thus the mid-prices reflect the output of "matrix pricing", whereby individual bond prices are imputed from market movements. This can induce spurious correlation with ETFs, which are the easiest reference point for such pricing. In addition, HY bonds have significant idiosyncratic risk, which increases the importance of actual trades for determining a reference point for pricing, as opposed to relying on matrix pricing. Due to these factors, the correlation between HY PTs and HY ETFs is not a useful measure of the ability to price and hedge with ETFs.

Instead, we calculate two (related) alternative measures of the "closeness" of a portfolio with ETFs, both based on ETF ownership of the individual line-items. First, we compute the average ETF ownership of all the bonds in each HY portfolio (*% ETF Ownership*). On average, 5.2% of the outstanding for bonds included in HY PTs (*Panel B* of **Table XI**, ***Panel B***). Second, we compute the percentage of the portfolio that is in bottom quartile of ETF ownership across the universe of HY bonds (*% Low ETF*). On average, 15% of a portfolio is comprised of these bonds. These measures are informed by the finding that the line-item PT benefit is lower for bonds that are not heavily owned by ETFs. TRACE data reveals that HY bonds in the bottom quartile of ETF ownership have significantly fewer trades and lower volumes than bonds that are more heavily owned by ETFs, after adjusting for liquidity (for a bond with median LCS, these non-ETF bonds have 20% fewer trades 17% lower volumes, shown in *Panel C* of **Table XI**).

The additional trading activity associated with ETF bonds is necessary to activate both ETF channels. Market-makers cannot map ETF prices into the price of a portfolio of bonds with few or no recent trades. This is particularly important in HY, due to the elevated idiosyncratic risk of HY bonds. Further, additional transactions in these bonds provide market-makers with more outlets for inventory, akin to the C/R results in IG (the additional trades and volumes for HY bonds heavily owned by ETFs far exceed any flows associated with the HY C/R process). As a result, greater ETF ownership and/or a lower percentage of low ETF ownership should be associated with a larger portfolio-level PT benefit.

Both measures are statistically significant and economically meaningful. An inter-quarter shift in *% ETF Ownership* results in a 12% improvement in the PT benefit. Similarly, an interquartile shift in *% Low ETF* results in a 9% reduction in the PT benefit (**Table XII**).

6. Alternative Explanations for PT Effectiveness

By focusing on the difference in execution costs between PTs and SSTs we control for common factors that drive execution costs, such as inventory and hedging costs (e.g., Goldstein and Hotchkiss (2020)), price transparency (e.g., Edwards, Harris, & Piwowar (2007)), and volatility, among others. We consider and reject five alternative explanations for the effectiveness of PTs: diversification, adverse selection and signaling, misallocation of the PT price across securities, investors swapping portfolios, and competition for market share.

Diversification

Li et al. (2023) apply our algorithm to study PTs and attribute the reduction in transaction costs to diversification rather than the ETF ecosystem. The authors argue that PTs diversify market-maker inventories, reducing idiosyncratic risk.

There are several issues with this explanation. First, market-makers operate large books with inventory measured in billions of dollars of gross longs and shorts spread over thousands of individual bonds. The practical scope to diversify this inventory through a portfolio trade (with a typical notional of c.$70mm) is quite limited, particularly since PT execution is not limited to smaller market-makers who may hold less inventory. Second, in support the diversification argument Li et al. (2023) show that the absolute reduction in transaction costs is greater (in basis-points) for riskier bonds, such as those with lower credit ratings. Such a comparison is flawed because, as we show in ***Table V***, lower rated bonds have substantially higher absolute transaction costs. When the benefits of PTs are appropriately normalized by the average transaction cost in their respective asset class, the percentage-point improvement is identical across IG and HY.

Finally, we examine long-short portfolios, which are an ideal laboratory to distinguish between the ETF and diversification hypotheses. Where long-only or short-only PTs contain mostly systematic risk, long-short PTs contain mostly idiosyncratic risk, particularly if the notional of the two legs is similar. Hence, if the benefits of PTs accrued mostly through a diversification channel, long-short PTs will not exhibit the same transaction cost benefits: executing a long-short PT adds to the idiosycratic risk in a market-maker's book. In constrast, the ETF channel applies equally to uni-directional and long-short PTs; the latter benefit if both legs can be individually priced and hedged using ETFs. This approach is better suited to distinguishing between diversification and the ETF channel than using the number of line items (as in Li et al) because larger portfolios are naturally more correlated with ETFs.

We first note that, over time, the percentage of long-short PTs has grown: from 38% of PT volumes in 2018 to 68% towards the end of 2023. This shows that market-makers have taken on more idiosyncratic risk over time through PTs, not less, and would be a surprising result if market-makers were worried about diversification (***Figure 6***).

Next, we formally compare the ETF and diversification channels. First, we compare long-short and unidirectional portfolios, by running a version of ***Model 4*** which combines a dummy for long-short portfolios with all of the relevant measures of the ETF channels; we run it separately for IG and HY PTs. IG long-short portfolios have slightly lower transaction costs than long-only or short-only portfolios (column (1),**Table *XIII***). In HY, long-short portfolios are not signicantly different from uni-directional portfolios (columns (3) and (5), **Table *XIII***). These results favor the ETF channel.

Second, we focus on the subsample of long-short portfolios and compare "balanced" and "unbalanced" portfolios. Within long-short portfolios, there is substantial variation in the relative notionals of the long and short leg, which we demonstrate by computing the weight (in terms of notional) of the short leg of each long-short portfolio[20] (***Figure 7*** contains the histogram of this weight). Approximately 30% of the long-short portfolios are balanced, defined as having a weight of the short leg between 45% and 55%. These contain mostly idiosyncratic risk since the systematic risk in two legs roughly nets. In contrast, unbalanced long-short portfolios contain more systematic risk.

Following this logic, we define a dummy variable *L-S Dummy Balanced* which equals one for balanced long-short portfolios, and zero for unbalanced portfolios. We restrict the sample to long-short PTs only, and regress $EHS - \widehat{EHS}$ on the relevant measures of ETF "closeness" in IG and HY (*Corr ETF* and *% Rightway ETF C/R* for IG and *% ETF Ownership* and *% Low ETF* in HY) and include *L-S Dummy Balanced* in the same specification (columns (2), (4), and (6) **Table *XIII***). There is no difference between balanced and unbalanced long-short portfolio in either IG and HY, in direct contrast to the

[20] By construction, we obtain the mirror image if we instead we looked at the long leg. We have tested specifications where we trimmed extreme values (close to 0 and close 1) and obtain similar results.

diversification channel. However, the ETF closeness variables remain economically and statistically relevant, in keeping with the ETF channel.

Finally, we note that long-short portfolios, particularly in IG, are larger on average than one-directional portfolios: the average long-short IG PT contains 109 line items, compared to only 79 line items for uni-directional PTs. In the Appendix we demonstrate that relationship between line items and execution is driven entirely by the larger size, and better execution, of long-short portfolios; once we control for the type of portfolio, the number of line items has no effect on the benefits of portfolio trading (**Table *A3.3***).

Signalling and adverse selection

Another possible explanation for the efficiency of PTs is that they subject market-makers to less risk of adverse selection. When trading an individual security, market-makers must account for the possibility that the investor (or market-maker) it is trading with has some private information. This is less of a concern at the portfolio level. It is more difficult to accumulate private information about a large basket of bonds, and investors are unlikely to execute many uninformed trades to provide "cover" for a small number of informed trades. If true, investors could even use the PT protocol to signal the absence of private information.

We perform two direct tests of the signalling hypothesis. First, we note that the index rebalancing trades utilized above in our causality tests are not polluted by private information. Rebalancing drives exogenous demand to trade around month-end from passive investors seeking to limit their tracking error that is divorced from any assessment of security value. Therefore, the analysis in ***Table VI***, where we rerun ***Model 1*** on this subset of trades, is a comparison of PTs and SSTs where neither sample is polluted by adverse selection. In IG, the coefficient on the PT dummy for rebalancing trades is very similar to that in the full sample, as is the percentage point reduction in EHS associated with PTs. In HY, the coefficient on the PT dummy is negative and significant, albeit slightly smaller in magnitude than in the full sample (-8.7bp versus -11.2bp). However, HY index rebalancing trades are in general cheaper than other trades, such that the percentage reduction in EHS is the equal to that in the full sample. Together, these results demonstrate that PTs remain efficient vis-à-vis SSTs even when signalling is not a concern.

Secondly, we examine the opposite end of the signalling spectrum and consider transactions which are very likely polluted by private information: those where an investor buys (sells) a bond and the last five trades in that bond were also buys (sells). The one-sided nature of the recent flows suggests that market-makers should be wary of taking an opposing position in that bond. Importantly, the market-maker can observe past trades in TRACE and its quote will reflect potential adverse selection regardless of whether the investor it is transacting with has that information (the market-maker is exposed regardless). Therefore, under the signalling hypothesis we expect little to no PT benefit for these trades.

For tractability, we test this hypothesis by combining the two pressure variables used in ***Model 1*** ($Buy\ Pressure_{i,j,t}$ and $Sell\ Pressure_{i,j,t}$) into one dummy variable ($Pressure_{i,j,t}$) that is equal to 1 if the trade in question is in the same direction as the previous five trades, and 0 otherwise. We then include that variable and its interaction with the PT dummy in an augmented version of ***Model 1***. The results are reported in **Table *XIV***. As in our baseline results, the coefficient on $PT_{i,j,t}$ is negative and significant; PTs have lower EHS than SSTs. Further, the coefficient on $Pressure_{i,j,t}$ is positive and significant in both IG and HY (2.5bp and 7.1bp respectively). As expected, the trades in question are more expensive to execute, by 17% in IG and 27% in HY, compared to the respective average EHS. This reflects the potential for adverse selection; the bid-offer is adjusted upwards when market-makers are concerned about their counterparty may have private information. However, the coefficient on the interaction term ($PT_{i,j,t}*\ Pressure_{i,j,t}$) is negative and significant. This is the opposite of what we would expect under the signalling hypothesis, which would predict a positive interaction term, undoing (at least some of) the PT benefit associated with these trades. We conclude that PTs retain their effectiveness even when market-makers are aware of the potential for adverse selection in the bond.

Taken together, these results provide strong evidence against the signalling hypothesis; if signalling were responsible for the improvement in EHS associated with PTs, one or both tests should show a much-reduced PT benefit. In fact, the evidence is almost too strong: market-makers are exposed to adverse selection, and it is unlikely that an investor in possession of private information would action it via a PT. Yet our results show that the relative cost of PTs is invariant to the potential for signalling. This seeming contradiction is explained by the large overlap between the line items in PTs and the bonds heavily owned by ETFs. ETFs increase price informativeness that reduces the potential for investors to acquire private information about the value of the securities they own (Choi, Kronlund and Oh (2022)). Put differently, illiquid bonds with limited ETF ownership have the greatest potential for adverse selection, as investors and market-makers have the least information about their prices. These bonds should be overrepresented in PTs if investors were using the protocol to signal an absence of private information. In practice, these bonds are significantly underrepresented in PTs. The underrepresentation of bonds that benefit from signalling but interfere with the ETF channels we identify provides additional support for our hypothesis.

Misallocation of the PT price

Another alternative explanation is that market-makers misallocate the PT price across line items in a manner designed to create a spurious (measured) reduction in the transaction cost. Such strategic misallocation is necessarily zero-sum. If some securities are priced artificially well, then some other securities must be priced artificially poorly, since the overall PT price is fixed. Our earlier results demonstrated that illiquid bonds in particular benefit from PT execution. If these gains are spurious, then market-makers must be mispricing some other set of bonds in the opposite direction. For obvious reasons, liquid bonds are unlikely to be mispriced. First, our analysis shows that they also benefit from PT execution, albeit to a lesser extent than illiquid bonds. Second, by virtue of their liquidity, investors have the greatest clarity on their value and will recognize any material mispricing.

However, most PTs contain bonds that trade in the PT protocol on a given day but not in the SST protocol. These "PT-only" bonds have no effect on our baseline specification because it relies on bonds that trade in both protocols on the same day. Further, these bonds tend to be illiquid, and thus potentially have stale mid prices. Market-makers could conceivably price these bonds from their true mid-price, but close to the (noisy) estimate of mid that is based on stale trading data, and then price bonds which have traded recently (and which drive our results) close to their (observable) mid, resulting in a spurious PT benefit.

The magnitude of the PT benefit and the paucity of PT-only bonds suggest that this is an unlikely explanation. PT-only bonds made up only about 20% of the average PT. These would need to be severely mispriced to explain the aggregate PT benefit: by 30bp in IG and 55bp in HY, assuming that all of the PT-only bonds had stale mid prices in the correct direction (such that they could be exploited to create a spurious benefit), and by even more if only a subset of these bonds did. We confirm that this type of misallocation is not responsible for our results with two formal tests. First, we examine the "drift" of the Bloomberg mid-price of the PT-only bonds after execution. If the PT benefit is spurious due to stale prices of PT-only bonds, then their mid-price should move against the investor over time, as the stated mid prices converge to their "true" value. However, in neither IG nor HY is there any significant adverse trend in mid-prices of these bonds in the 15 trading after PT execution

(**Figure *8***). Second, we regress the portfolio-level PT benefit against the proportion of the portfolio that was PT-only on the day of execution (**Table *XV***). A larger proportion of PT-only bonds should be associated with a larger PT benefit under the misallocation hypothesis, as it provides market-makers more leeway to exploit stale mid prices. In HY the coefficient on the percentage of PT-only bonds is insignificant both statistically and economically. In IG, the coefficient is negative and marginally statistically significant. However, the intercept (-5bp) is 80% of the average PT benefit; portfolios with no PT-only bonds retain nearly all the benefit of PT execution. Further, for the average PT (with 20% of bonds that traded only in PTs), the coefficient implies only a 0.3bp improvement in the portfolio-level PT benefit, which is a full order of magnitude below the measured PT benefit.

Finally, the overlap of PTs and ETFs is additional evidence against the misallocation hypothesis. Bonds with limited ETF ownership are those with the most potential for misallocation, and market-makers should be happy to include them in PTs if there were looking for opportunities to misprice bonds. Yet they are rarely included in PTs. We conclude that the PT benefit we measure is not driven by a misallocation of the PT price.

Investors swapping portfolios

It is possible that investors "swap" PTs with other investors, who want to trade the same bonds but in the opposite direction. If the market-maker acts as an agent, lining up both sides of the PT, the bid-offer would be much lower than the average bid-offer of SSTs, which include many principal trades. However, we find that that offsetting PTs occur in less than 0.05% of the PTs in our sample, which speaks strongly against the theory.

Competition for market share

Market-makers might use portfolio trades to gain market share, and the competition to win these trades could drive the transaction cost they charge. 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 *9* we plot the time series of the difference between the average *EHS* of transactions executed via the PT and SST protocol over the period Jan 1st 2020 – December 31st 2021. We also show PT volumes as a percentage of total TRACE 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 market-makers were "buying" market share, as they would have increased the benefit of portfolio trading to entice more participation.

PT clustering

A final concern is that PTs absorb "easy" liquidity, leaving more difficult trades to the SST protocol. For example, investors might bundle easy (and low cost) trades into PTs, which would lead us to falsely measure a PT benefit. Although both the use of bond-day fixed effects and the index rebalancing analysis both address this concern, we compare the average *EHS* of SSTs on days with high and low PT activity as an additional check. If PTs merely absorb (at least some of) the "easy" trades available to investors, then the average *EHS* of SSTs would be lower on days with relatively few PTs. Daily PT volumes range from 5.7% of TRACE (P25) to 10.9% of TRACE (P75) (Panel A (IG) and Panel B (HY), **Table *XVI***). However, the average SST *EHS* is the same on days with low and high PT activity. This also confirms that PTs are not clustered on days with high liquidity.

7. Discussion: The Uses and Limitations of PTs

Our analysis of portfolio trades explains why they are effective and allows us to contextualize their potential use and growth vis-à-vis SSTs. Corporate bond liquidity has historically been challenged by an onerous matching problem. The large number of CUSIPs makes it difficult for market-makers to find an immediate and exact offsetting trade for individual bond positions they take into inventory. Individual bonds are difficult to hedge, and the risk associated with inventory management is a drag on the provision of immediacy. This explains why corporate bond liquidity has historically been challenging and expensive.

Portfolio trades are an elegant solution to this matching problem. The deep secondary market liquidity of ETFs allows market-makers to efficiently price and hedge portfolios, and market-makers can more easily offload the inventory they accumulate through PTs so long as the bonds are heavily owned by ETFs, either through traditional investor flows or the creation and redemption process. Reductions in the costs and risks borne by market-makers are reflected in the cost of portfolio trades, just as they are in other trading protocols, such as agency trading. In fact, the material savings associated with PTs demonstrates just how effectively these channels mitigate the challenges associated with intermediating bond trades.

The decline in transaction costs is a major draw for investors, and unsurprisingly the use of PTs has grown rapidly since 2018 (**Table *II*** Panel A (IG) and Panel B (HY)). In the IG market, the number of PTs grew by a factor of 2.5 between 2018 (1,950 unique PTs) and 2021 (over 4,900). Portfolio trading volume increased from $81 billion in 2018 to $311 billion in 2021. In percentage terms, PTs comprised 7% of trading volume in 2021, off a base of c.1% in 2018. This growth continued thereafter; in 2024 we estimate that PTs comprised c.14% of volume in the IG market. We observe similar patterns in the HY market, where PTs also comprised 7% of volumes in 2021, which grew to 10% in 2024. This rapid growth demonstrates that the protocol has been quickly adopted by many market participants.

While this growth is impressive, most trades are still executed using the SST protocol. In fact, despite their low transaction costs, we do not expect PTs to overtake or replace SSTs. Portfolio trading is best interpreted as a complement to SSTs, and investors will toggle between the two protocols depending on their investment needs and the market environment. Where they have trading needs that are well-suited to portfolio trading, they will use it, to take advantage of the reduction in transaction costs. But SSTs will be preferred for the many trading needs that are poorly suited to PTs.

The major constraint on the use of PTs comes from the source of their efficiency: the linkages to ETFs. Careful portfolio construction is essential to activate the ETF channels that we identify and thus unlock the reduction in transaction costs. Investors cannot simply bundle a random selection of bonds into a PT and expect to transact at a low bid-offer. In fact, our analysis suggests there are two important constraints on portfolio construction; together, these define the use cases for which PTs are appropriate.

The first constraint is that PTs must include a sufficient number of liquid bonds, even though the benefits of portfolio trading accrue disproportionately to illiquid bonds. The liquid bonds allow investors to "crowd-source" liquidity for illiquid bonds, by enabling market-makers to utilize the information from the ETF market. The liquid bonds provide a reference point for the price of the portfolio, which market-makers can adjust using ETF prices. In contrast, market-makers would struggle to translate the real-time secondary price of an ETF into the price of a portfolio comprised mostly or entirely of bonds that never trade.

As evidence of this constraint, we note that PTs are not exclusively comprised of illiquid bonds, despite being less liquid on average than the index. In aggregate, the protocol is utilized for bonds across the liquidity spectrum (***Figure 3***). This is also true of individual PTs. We define Illiquid (Liquid) PTs as those with higher (lower) notional-weighted LCS than the trade volume-weighted LCS of the bonds in the BBG IG Index. In **Figure *10***we plot the notional-weighted percentage of bonds in the two most liquid LCS quintiles for each category; even illiquid PTs contain a substantial proportion (37%) of liquid bonds.

The second constraint is that the portfolio must contain the right illiquid bonds: those heavily owned by ETFs. These bonds are easier for market-makers to offload, both through increased investor flows and through the C/R process. Illiquid bonds that are not heavily owned by ETFs are more difficult to offload and thus reduce PT efficiency. This constraint is binding: PTs over-index for illiquid bonds, but only those heavily owned by ETFs.

Our results indicate that both channels must be activated to reduce transaction costs. Inventory in illiquid ETF bonds are easier for market-makers to offload, regardless of whether the risk was sourced via PTs or SSTs. But transactions in these bonds are cheaper when they are executed as part of a PT because the crowd sourcing allows for efficient pricing and hedging. Similarly, the effectiveness of PTs for liquid bonds is limited; the pricing and hedging capabilities afforded by ETFs apply to portfolios comprised entirely of liquid bonds, but due to their already-high liquidity, execution via PTs has a smaller effect on their transaction costs. The real gains come from including illiquid bonds owned by ETFs.

Trading needs that activate both channels can be implemented using PTs; those that fail to activate one or both will be executed as SSTs. Ensuring that a PT can be priced and hedged with ETFs limits portfolio trading to the subset of investment objectives that involve trading a large and diversified portfolio of bonds. Credit investors typically own portfolios with hundreds or even thousands of CUSIPs. As a result, they naturally execute many transactions that meet this criterion. An obvious example is a pure “beta” play, where an investor trades an index-like portfolio. These are easy for market-makers to hedge with ETFs (the largest ETFs are themselves beta plays) and can be used to satisfy many liquidity needs, such as daily inflows and outflows experienced by mutual funds.

But investors are not limited to index-like portfolios. 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 (**Table *XVII***). We classify a PT as having a maturity, sector, or credit rating motivation if its HHI along that dimension is at least 50% higher than the respective HHI of the Index. A significant fraction of portfolios is concentrated along at least one of these dimensions, with the maturity-type strategy being the most common (35% of IG portfolios, typically focused on longer-dated bonds). This indicates that investors use

PTs to manage key macro characteristics of their overall portfolios. Importantly, these PTs meet the requirements necessary to benefit from efficient execution. They have the requisite mix of liquid and illiquid bonds (e.g., 80% of the maturity-based PTs have aggregate portfolio liquidity below that of the broad index) and can be priced and hedged using ETFs. In addition to being highly correlated to the broad ETFs, they can be hedged using the increasing array of narrow ETFs (in IG, the largest narrow ETFs are tenor specific, and in HY both rating- and sector-specific ETFs have emerged).

The prominence of long-short PTs is additional evidence that investors use PTs to fine-tune their exposures (long-short PTs are definitionally not pure beta plays). Index rebalancing is one use case for long-short PTs, as it requires both buying and selling many individual CUSIPs. Another example is systematic trading of corporate bonds. The high bid-offer typical of SSTs has long deterred this type of investing, as it requires frequent transactions in many CUSIPs. Conversations with bond traders suggest that systematic investing has become more important in both the IG and HY markets since the advent of portfolio trading, motivated in at least some cases by the changes in the liquidity landscape.[21] This development is an interesting avenue for future work, as previous research has documented the effect of systematic investing on the distribution of returns in other asset classes, and it may have similar effects on corporate bonds.

However, many trading strategies run afoul of these constraints, including transactions in specific securities, issuers, or narrow market segments (e.g., large banks). These are difficult to hedge using ETFs and must be executed as an SST. There are many potential catalysts for these types of trades, such as earnings beats and misses, changes in capital structure policy, and changes in regulation affecting specific issuers or market segments, and they are common trading needs for fundamental credit investors.

The emphasis on bonds that are heavily owned by ETFs similarly limits the applicability of PTs for a large portion of the investable universe. For example, LQD owns only 30% of the bonds in the Bloomberg IG Corporate Bond Index. While the number of bond ETFs continues to grow, and with it the average ETF ownership, a significant fraction of bonds is unlikely to be owned in sufficient size by ETFs to feature heavily in portfolio trading. These bonds will continue to trade primarily in the SST protocol. This also implies that any investor who owns bonds across the spectrum of ETF ownership cannot use PTs to trade pro-rata

[21] See *State-Street-Global-Advisors-and-Barclays-Quantitative-Portfolio-Strategy-Announce-Collaboration-on-Active-Systematic-Fixed-Income-Strategies*.

slices of its portfolio, which could limit their use. For example, while a mutual fund could use PTs to manage small daily inflows and outflows, it would struggle to use PTs to satisfy sustained flows in one direction, as the limitation to ETF bonds would eventually distort its holdings. As discussed above, this limitation also prevents investors from using PTs to signal that they do not have private information about bond valuations. This conceptually appealing use of PTs to overcome the lemons problem faced by market-makers would be most useful for the very bonds that are rarely transacted using the protocol.

Our analysis of the COVID period also indicates the potential for significant time-series variation in the effectiveness (and availability) of PTs. Sustained one-sided flows interfere with the ETF C/R channel; if investors are all trying to buy (sell) while the ETFs are in create (redeem) mode, then market-makers lose the ability to offload inventory to the ETFs. Given that both channels are essential to PT efficiency, sustained swings eliminate the benefit of PTs for most investors (i.e., those going the same direction as the market). Although we have no episodes to point to, a disruption in ETF trading would have a similar effect, as it would jam the pricing and hedging channel.

Potential disruptions in the availability or efficacy of PTs raise concerns about their effect on systemic risk in the bond market that are interesting for both regulators and researchers. Of course, investors can always transact via SSTs, but there are several reasons why this is insufficient to dismiss these concerns. First, PTs may affect the competitive dynamics for market-makers. For example, they may increase the market share of the largest market-makers, as PTs require large gross balance sheet commitment and the required investments in technology to effectively trade and risk manage PTs may be large. With fewer or more concentrated market-makers, liquidity may suffer in periods of stress. Alternatively, the rise of PTs could increase the market share of market-makers that specialize in PTs, ETFs, and algorithmic trading, which may be more likely to withdraw from the market in times of stress. Second, investment styles (such as systematic credit) whose viability is based on PTs could experience fire sales if the protocol was disrupted for a sustained period.

Finally, improvements in bond liquidity may reduce the demand for public credit from investors with limited trading needs. The corporate bond liquidity premium is an important source of returns for these investors, and it could decline as innovations like PTs improve liquidity. This may contribute to the demand for private credit, which has grown alongside portfolio trading. Exploring the potential connections between improving liquidity in public markets and developments in private markets is another avenue for future work.

List of Figures

Figure 1: Client Inquiries

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

Alt text: The figure shows an increase in both the count and volume of PT inquiries over time.

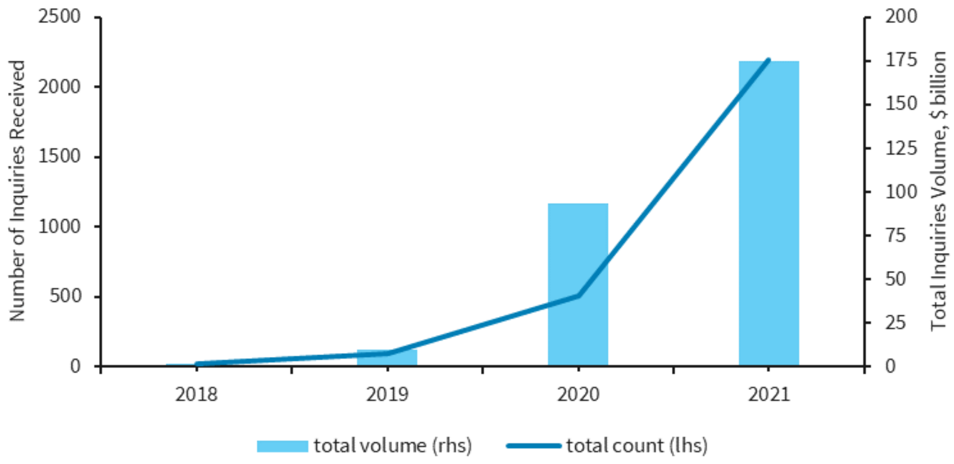

Figure 2: Flowchart of the Methodology Process

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

Alt text: The figure summarizes the four step algorithm we use to identify PTs in the TRACE dataset.

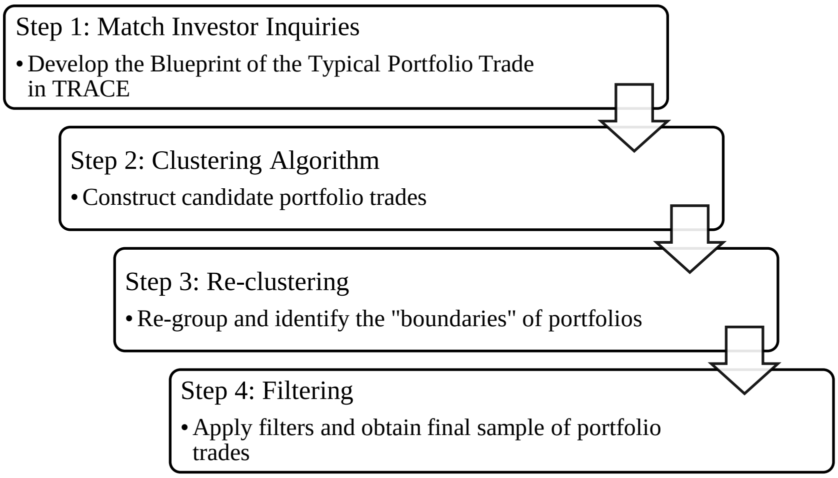

Figure 3: Distribution of Portfolio Volume by Liquidity Quintile in IG Portfolios

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.

Alt text: The figure shows that PTs over-index for less liquid bonds vis-à-vis overall volumes in TRACE.

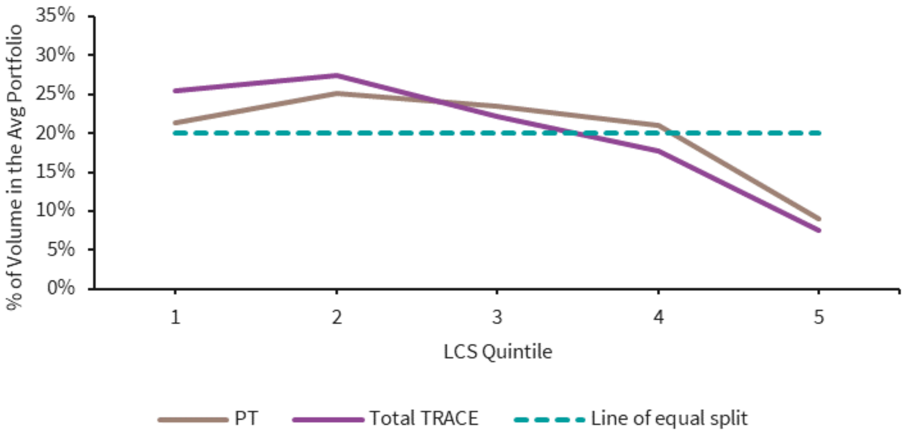

Figure 4: Overlap Between Portfolio Trades and ETFs

The figure shows the overlap between the line items portfolio trades and the monthly holdings of ETFs. Panel A shows the overlap with the largest ETFs (LQD, VCIT, VCSH, IGIB, IGSB in IG, and HYG, JNK, SHYG and SJNK in HY), Panel B shows the overlap of IG portfolios with the largest IG ETF (LQD) and Panel C shows the overlap of HY portfolios with the largest HY ETF (HYG). Data based on portfolio trades identified by our ML algorithm executed during the period 1st January 2021 – 31st December 2021.

Panel A: Overlap with the largest ETFs

Alt text: The figure shows that a histogram of the proportion of line items in IG PTs that are owned by IG ETFs, which peaks at around 90%.

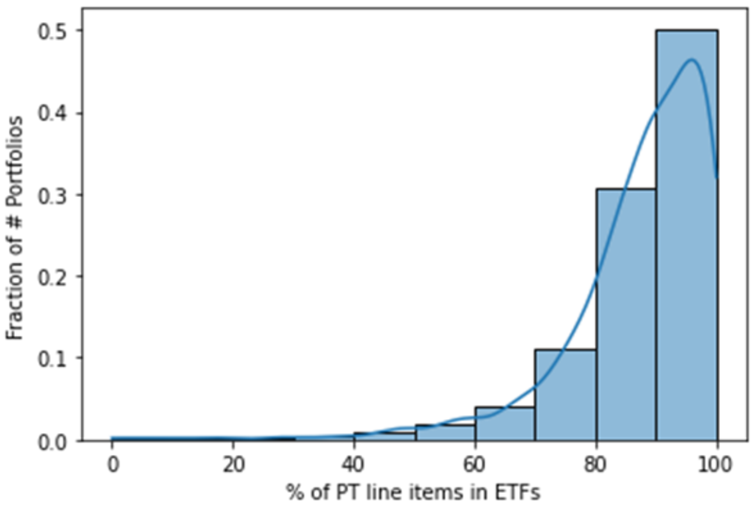

Panel B: Overlap with LQD (IG portfolios)

Alt text: The figure shows the same histogram but for LQD only, one of the largest IG ETFs; it peaks at 60%.

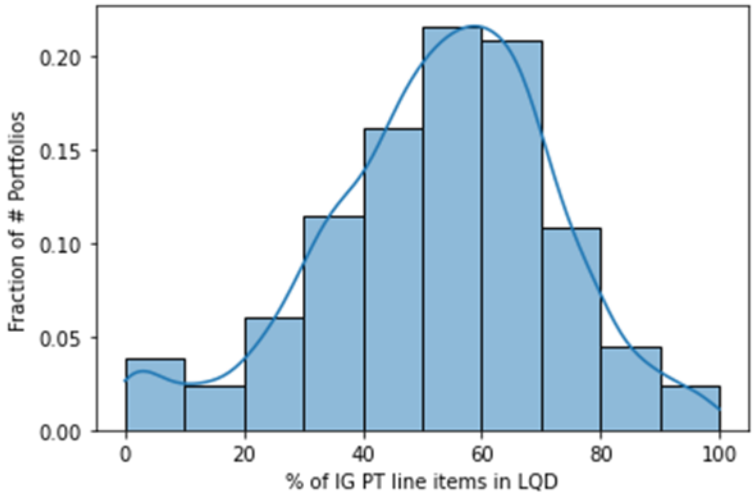

Panel C: Overlap with HYG (HY portfolios)

Alt text: The figure shows the same histogram but for the overlap between HY PTs and HYG, a large HY ETF; it peaks at 80%.

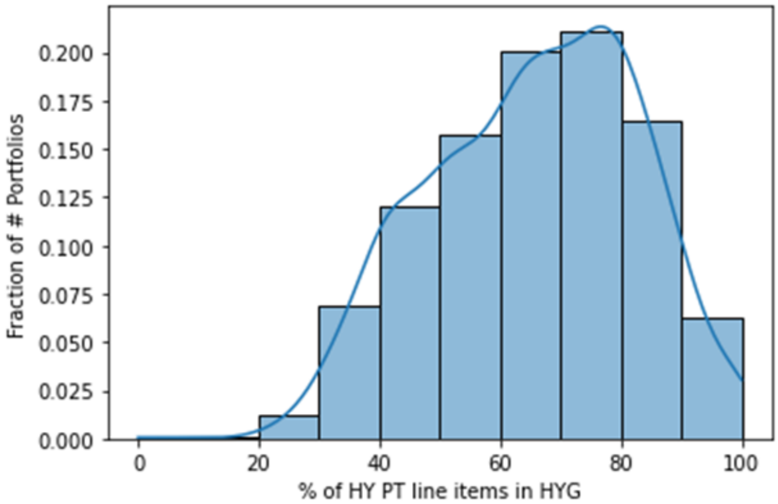

Figure 5: **Transaction Cost Analysis – Identification Strategy**

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

Alt text: The figure shows that over 20% of bond-days trade in both the SST and PT protocol in our sample.

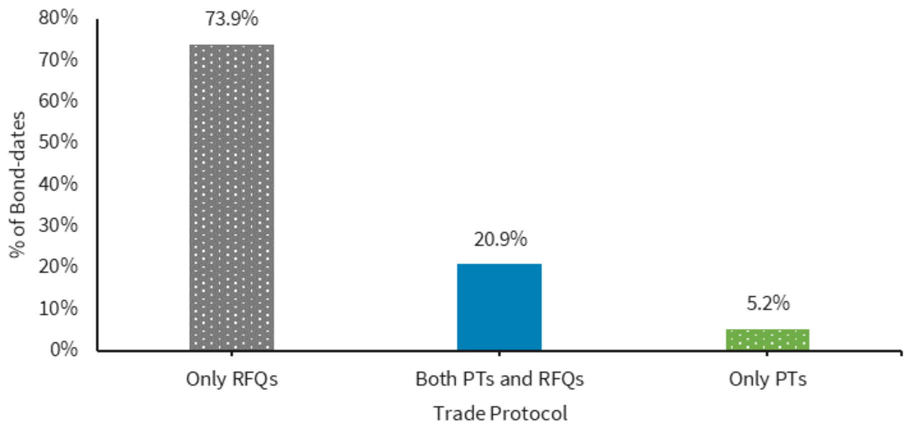

Figure 6: **Long-short Portfolio Trades over Time**

The figure shows the percentage of portfolio volume in long-short portfolio trades.

Alt text: The figure demonstrates that long-short PTs have increased in frequency over time, from 37% of PTs in 2018 to 68% in 2023.

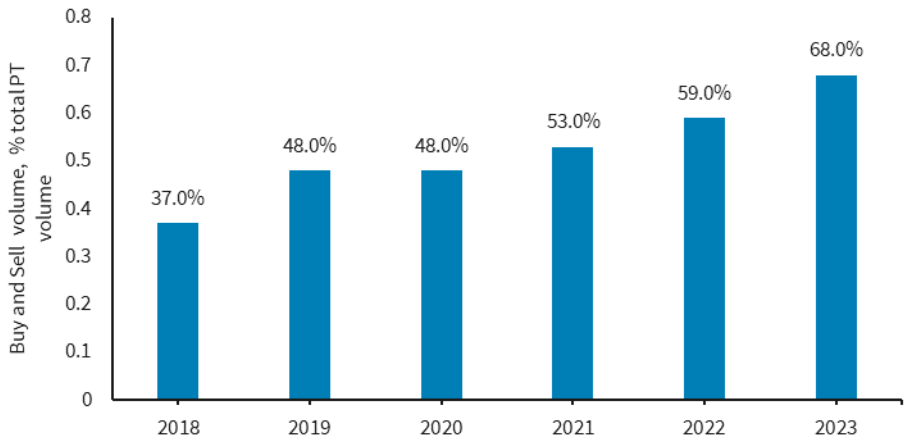

Figure 7: Balanced vs. Unbalanced Long-short Portfolios

The figure shows the distribution of the percentage of total portfolio notional in the short leg.

Alt text: The figure shows that long-short PTs include both balanced (i.e., the long and short legs are approximately equal notional) and unbalanced PTs.

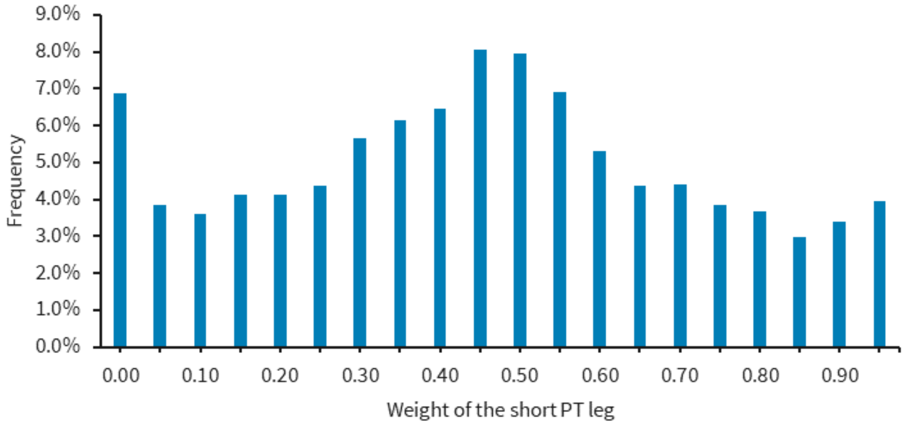

Figure 8: Drift of PT-only Bonds After PT Execution Date

The figure shows the notional-weighted average drift of bonds that only trade in PT form for the 15 days after PT execution, measured in bp. Negative values indicate the bonds moved in the investor's favor, and positive values indicate the bonds moved against the investor. The dashed lines indicate the average drift in each market necessary to explain the PT benefit measured for bonds that traded in both protocols on the day of execution. Data are based on IG and HY portfolio trades identified by our ML algorithm executed during the period 1st January 2020 – 31st December 2021.

Alt text: The figure shows that there is no net drift of the marks of bonds included in PTs against the investor over the three weeks after execution.

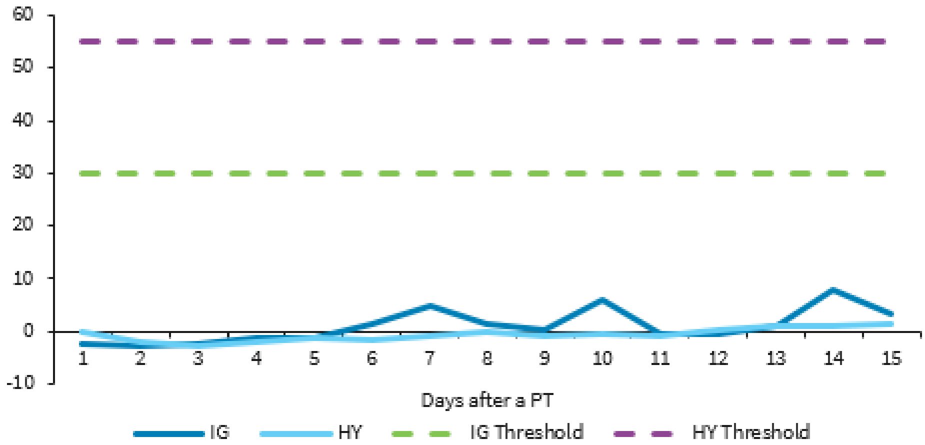

Figure 9: Portfolio Trade Execution Costs Over Time

The figure shows the difference between the average EHS of transactions executed in the PT and SST 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 SST protocol. For more details on the definition of our TRACE universe refer to Section 2.3. Data are based on IG and HY portfolio trades identified by our ML algorithm executed during the period 1st January 2020 – 31st December 2021.

Alt text: The figure shows that the PT benefit does not vary with the proportion of trades done in PT form.

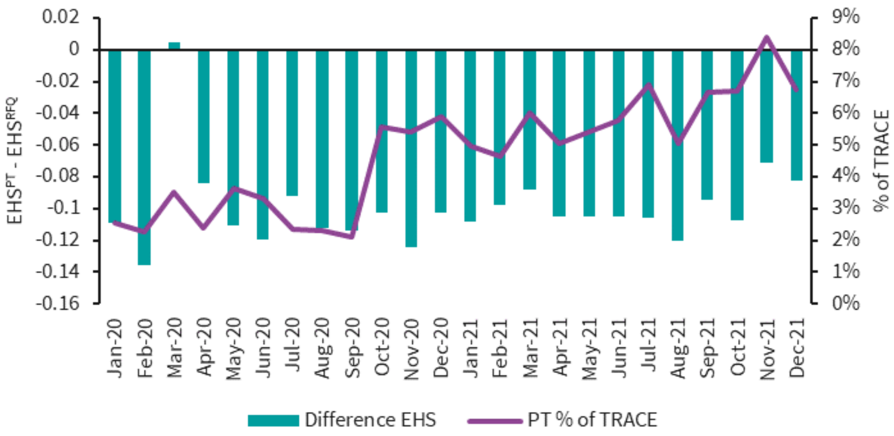

Figure 10: Mixing Liquid and Illiquid Bonds in IG Portfolios

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.

Alt text: The figure demonstrates that there even PTs that heavily feature illiquid bonds also include at least some liquid bonds.

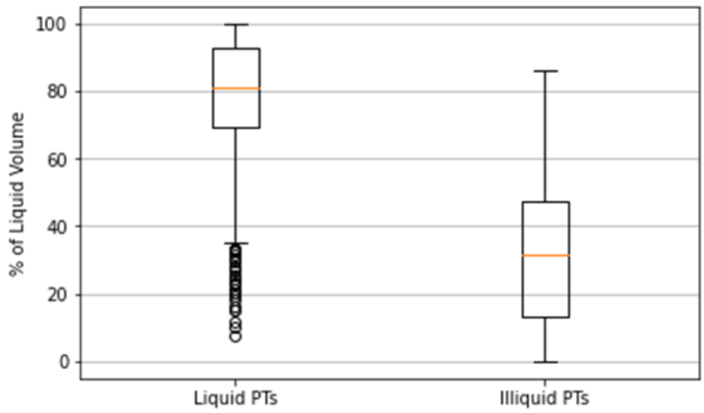

List of Tables

Table I: 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.

Alt text: The table shows sample PT inquiry, with the direction and size for each line item.

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.

Alt text: The table shows summary statistics for IG and HY PTs over time, in terms of volume and proportion of TRACE.

Panel A: IG

	# Bond-PT Obs.	**# of PTs**	**$ Volume (bln)**	**% of TRACE**
Panel A: Aggregate				
2018-2021	998,975	12,107	696	3.5%
Panel B: Time Series				
2018	107,541	1,950	81	1.1%
2019	175,224	2,265	127	1.7%
2020	245,774	2,978	177	3.1%
2021	470,436	4,914	311	6.9%

Panel B: HY

	# Bond-PT Obs.	**# of PTs**	**$ Volume (bln)**	**% of TRACE**
Panel A: Aggregate				
2018-2021	375,090	4,157	241	3.3%
Panel B: Time Series				
2018	41,363	592	23	1.3%
2019	47,471	578	31	1.8%
2020	89,508	927	61	3.2%
2021	196,748	2,054	127	6.9%

Table III: PT Summary Statistics

The table presents summary statistics of PTs. Panel A contains characteristics at the portfolio level, and Panel B contains characteristics of the bonds in the portfolios. Portfolio LCS, Maturity and Age (time since issuance) are computed as a notional-weighted average of the PT line items. The last row, *TRACE ex PT*, reports the volume-weighted LCS, maturity and bond age for all non-portfolio trades in TRACE. For validation purposes, we include statistics for the investor inquiries (INQ) that we could match in TRACE alongside those of the portfolio trades (PT) identified using our ML algorithm.

Alt text: The table shows summary statistics of the PTs we identify and the line items that are included in them.

	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 IV: 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.

Alt text: PT frequency increases for bonds that are less liquid and more heavily owned by ETFs, in both IG and HY.

Panel A: IG

		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%	-

Panel B: HY

		ETF Ownership					
		Low	2	3	4	High	H-L
Illiquidity (LCS)	Low	1.9%	6.7%	8.9%	10.2%	10.8%	8.9%
	2	3.9%	7.6%	9.4%	10.6%	11.4%	7.5%
	3	3.5%	10.4%	10.%	11.3%	12.4%	5.9%
	4	6.5%	8.8%	10.6%	10.9%	12.%	5.7%
	High	6.2%	6.5%	10.6%	11.9%	11.2%	5.0%
	H-L	4.3%	-0.2%	1.7%	1.7%	0.4%	-

Table V: Transaction Costs of Portfolio Trades

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

The table reports transaction-level regressions of effective half spread ($EHS_{i,j,t}$)on a portfolio trade dummy ($PT_{i,j,t}$) 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_{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}$) and a dummy equal to one if all five previous trades were investor buys. All continuous variables are winsorized at the 1% level. Regressions include a full set of bond-date fixed effects ($\boldsymbol{\lambda_{j,t}}$). Columns (1) and (2) estimate the model for IG portfolios, column (3) and (4) for HY portfolios. 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.

Alt text: The table shows that PTs have lower EHS than SSTs in the same bonds on the same day, in both IG and HY.

Panel A: Point Estimates				
	IG		HY	
	(1) TRACE Portfolios	(2) Investor Inquiries	(3) TRACE Portfolios	(4) Investor Inquiries
PT Dummy	-6.10*** (-24.98)	-5.32*** (-4.72)	-11.17*** (-27.94)	-8.20*** (-12.98)
Block Trade	-2.92*** (-15.28)	-2.62*** (-13.86)	-9.66*** (-43.73))	-8.93*** (-39.92)
Lag EHS	-0.16*** (-66.85)	-0.16*** (-66.93)	-0.04*** (-9.58)	-0.04*** (-9.59)
Morning Trade	-0.41*** (-3.92)	-0.10 (-1.03)	-4.02*** (-18.21)	-3.41*** (-17.79)
Sell Pressure	0.95*** (2.87)	0.98*** (3.02)	5.11*** (7.79)	5.21*** (7.91)
Buy Pressure	3.19*** (8.21)	3.30*** (8.48)	7.23*** (12.89)	7.39*** (13.05)
Bond-Date FE	YES	YES	YES	YES
Bond-trade Observations	4,868,961	4,868,961	1,650,802	1,650,802
Sample Period	Jan 1st 2021 – December 31st 2021			
Panel B: Percentage Improvement				
Mean EHS SSTs	14.5bp	13.9bp	26.7bp	25.5bp
Improvement vs. SSTs	42.1%	38.3%	41.7%	32.2%

Table VI: Causality – Index rebalancing

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

The table reports transaction-level regressions of effective half spread ($EHS_{i,j,t}$) on a portfolio trade dummy ($PT_{i,j,t}$) and a set of transaction-level controls (collected in vector $Z_{i,j,t}$: a dummy equal to one if a trade is larger than $5 million ($Block_{i,j,t}$), previous trade EHS ($Lag\ EHS_{i,j,t}$) and a dummy equal to one if the trade is executed before 13:00 EST). The sample is limited to investor buy (sell) trades in bonds which entered (exited) the Bloomberg Investment Grade Corporate Bond Index, executed during the three days before or after month-end. All continuous variables are winsorized at the 1% level. Regressions include a full set of bond-date fixed effects ($\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.

Alt text: The table shows that the baseline results from the prior table are unchanged after limiting the sample to trades that buy (sell) bonds that are added (dropped) from the major IG and HY corporate bond indices in a 3-day window around index rebalancing.

Panel A: Point Estimates				
	IG		HY	
	(1) TRACE Portfolios	(2) Investor Inquiries	(3) TRACE Portfolios	(4) Investor Inquiries
PT Dummy	-7.52***	-10.05***	-8.66***	-9.51***
	(-5.81)	(-2.82)	(-5.62)	(-7.22)
Block Trade	-6.09***	-5.44***	-7.63***	-1.17***
	(-5.05)	(-4.84)	(-5.70)	(5.63)
Lag EHS	-0.02	-0.02	0.05	0.05
	(-0.82)	(-0.81)	(0.94)	(0.96)
Morning Trade	-2.95**	-2.30**	-1.95	-0.64
	(-2.55)	(-1.96)	(-0.64)	(-0.21)
Bond-Date FE	YES	YES	YES	YES
Bond-trade Observations	25,413		12,216	
Sample Period	Jan 1st 2021 – December 31st 2021			
Panel B: Percentage Improvement				
Mean EHS SSTs	17.7bp	16.7bp	20.5bp	18.9bp
Improvement vs. SSTs	42.5%	60.2%	42.2%	50.3%

Table VII: Portfolio Trades By Illiquidity Profile and ETF Ownership

$$\boldsymbol{EHS_{i,j,t} = \beta_1 PT_{i,j,t} + \beta_2 PT_{i,j,t} \times Illiq\ Dummy_{j,t} + \beta_3 PT_{i,j,t} \times ETF\ Own_{j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t}}$$

The table reports transaction-level regressions of effective half spread on a portfolio trade dummy ($PT_{i,j,t}$), an interaction term of the PT dummy with an illiquidity dummy ($Illiq\ Dummy_{j,t}$), and an interaction term of the PT dummy with a dummy for low ETF ownership ($ETF\ Own_{j,t}$). $Illiq\ Dummy_{j,t}$ equals to one for bonds in the most illiquid quintile of LCS. $ETF\ Own_{j,t}$ equals one for bonds in the lowest quartile (half) of ETF ownership in the IG (HY) market. Panels (2) and (4) also contain a portfolio-level measure of the extent of low ETF ownership ($Port\ ETF_{i,j,t}$) and its interaction with the PT dummy. $Port\ ETF_{i,j,t}$ is equal to the notional-weighted proportion of bonds in the lowest quartile (half) of ETF ownership in each IG (HY) portfolio (and is equal to 0 for SSTs). 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_{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}$) and a dummy equal to one if all five previous trades were investor buys ($Buy\ Pressure_{i,j,t}$). All continuous variables are winsorized at the 1% level. All regressions include a full set of bond-date fixed effects ($\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.

Alt text: The table shows that in both IG and HY, less liquid bonds and bonds more heavily owned by ETFs have a greater PT benefit.

	IG		HY	
	(1)	(2)	(3)	(4)
PT Dummy	-5.48*** (-37.07)	-5.43*** (-20.03)	-10.15*** (-32.59)	-11.92*** (-12.66)
PT × Illiq Dummy	-3.14*** (-5.71)	-3.13*** (-5.72)	-6.35*** (-16.27)	-6.32*** (-16.27)
PT × ETF Own	2.38*** (11.10)	1.52*** (4.13)	1.15*** (4.08)	-4.06*** (-3.83)
Port ETF	--	-0.79 (-0.26)	--	3.66** (1.95)
PT × Port ETF	--	6.98** (2.33)	--	8.63*** (4.37)
Block Trade	-2.93*** (-17.42)	-2.93*** (-4.14)	-9.62*** (-47.53)	-9.62*** (-47.36)
Lag EHS	-0.157*** (-72.68)	-0.16*** (-72.69)	-0.03*** (-10.67)	-0.03*** (-10.67)
Morning Trade	-0.42*** (-4.13)	-0.42*** (-4.14)	-4.02*** (-19.75)	-4.05*** (-19.83)
Sell Pressure	0.94*** (2.985)	0.94** (2.99)	5.10*** (8.63)	5.10*** (8.64)
Buy Pressure	3.198*** (8.49)	3.20*** (8.49)	7.29*** (14.64)	7.28*** (14.63)
Bond-Date FE	YES	YES	YES	YES
Bond-trade Observations	4,868,961	4,868,961	1,650,802	1,650,802
Sample Period	Jan 1st 2021 – December 31st 2021			

Table VIII: Portfolio-level Characteristics (IG)

The table reports summary statistics of $EHS_p - \widehat{EHS}_p$ (Panel A) and of the portfolio-level characteristics (Panel B) used in Section 5: *Corr ETF* (past 30-day correlation between portfolio and LQD returns), *% ETF 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 ETF C/R* (% of line items that were "rightway" for the LQD C/R process on the day the portfolio was executed, and defined on days with significant ETF C/R activity), and *PT LCS*, the weighted average LCS of the line-items in the portfolio (measured in percentage).

Alt text: The table shows summary statistics of the relationship between IG PTs and IG ETFs.

	Mean	Std	P25	P50	P75
Panel A: Predicted Cost Benefit of Portfolio Trading					
$EHS_p - \widehat{EHS}_p$	-5.4bp	6.0bp	-9.0bp	-5.4bp	-2.1bp
Panel B: Characteristics Linked to the ETF Ecosystem					
Corr ETF	0.79	0.19	0.71	0.87	0.93
% ETF C/R	35.9%	16.2%	24.5%	35.5%	47.8%
% Rightway ETF C/R	10.2%	14.8%	0.0%	4.2%	13.8%
PT LCS	0.62	0.23	0.44	0.57	0.74

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

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

The table reports the results of cross-sectional regressions of actual (*EHS*) minus predicted ($\widehat{EHS}$) effective half spread of portfolio ***p*** on portfolio-level characteristics: *Corr ETF* (past 30-day correlation between portfolio and LQD or HYG returns), *% ETF 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 ETF C/R* (% of line items that were "rightway" for LQD on the day the portfolio trade was executed, and defined on days with significant ETF C/R activity). The portfolio-level weighted average LCS (*PT LCS*) is included as a control. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Alt text: The table shows that IG PTs that are more correlated to IG ETFs and have greater overlap with the C/R process have a greater PT benefit.

	$EHS_p - \widehat{EHS}_p$		
	ETF Hedging and Pricing	ETF C/R	
	(1)	(2)	(3)
Const	-4.03*** (-10.12)	-4.41*** (-15.82)	-4.67*** (-11.93)
Corr ETF	-1.03*** (-2.2)	-	-
% ETF C/R	-	-0.013*** (-2.35)	-
% Rightway ETF C/R	-	-	-0.054*** (-5.88)
PT LCS	-0.85*** (-2.23)	-0.008*** (-2.14)	-0.003) (-0.55)
Sample	4,914 PTs		1,987 PTs

Table X: ETF C/R: COVID-19 Sample

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

The table reports transaction-level regressions of effective half spread ($EHS_{i,j,t}$) on a portfolio trade dummy ($PT_{i,j,t}$) and a set of controls. The model is estimated on data for IG corporate bonds during the period March 10, 2020 to March 31st, 2020. Column (2) includes $Against\ Market_{i,j,t}$, which is a dummy equal to 1 when trade *i* went against the prevailing flows (i.e., was a buy between March 10 and March 19, and a sell from March 20 to March 31), and its interaction with the PT dummy. 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_{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}$) and a dummy equal to one if all five previous trades were investor buys ($Buy\ Pressure_{i,j,t}$). All continuous variables are winsorized at the 1% level. Regressions include a full set of bond-date fixed effects ($\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.

Alt text: The table shows that most PTs did not obtain better execution than SSTs during March 2020, although those PTs that were against the prevailing market flows did have lower transaction costs than similar SSTs.

	(1)	(2)
PT	17.14 (1.41)	39.19** (2.07)
Against Market	--	11.35 (0.57)
PT × Against Market	--	-53.25** (2.03)
Block Trade	-41.51*** (-13.99)	-41.83*** (-12.90)
Lag EHS	-0.15*** (-18.71)	-0.15*** (-18.63)
Morning Trade	-7.00*** (2.40)	-6.87*** (-2.40)
Sell Pressure	26.82*** (5.18)	26.76*** (5.00)
Buy Pressure	-30.83*** (-3.74)	-30.15*** (-3.76)
Bond-Date FE	YES	YES
Bond-trade Observations	275,018	275,018
Sample Period	March 10st 2020 – March 31st 2020	

Table XI: Portfolio-level Characteristics (HY)

The table reports summary statistics of $EHS_p - \widehat{EHS_p}$ (Panel A) and of the portfolio-level characteristics (Panel B) used in Section 5: *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of HY ETFs), *% Low ETF* (% of line items in the bottom quartile of ETF ownership), and *PT LCS* (the weighted average of the LCS of the line-items in the portfolio, measured in percentage). Panel C includes the regression of trade count and trade volumes for TRACE data from 2021, regressed against bond LCS (averaged over the year) and a dummy variable *Low ETF* equal to 1 for bonds in the lowest quartile of ETF ownership. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Alt text: The table shows summary statistics of the relationship between HY PTs and HY ETFs (Panels A and B) and that low ETF ownership is correlated with fewer trades and lower overall trade volumes.

	Mean	Std	P25	P50	P75
Panel A: Predicted Cost Benefit of Portfolio Trading					
$EHS_p - \widehat{EHS_p}$	-10.2bp	6.2bp	-14.7bp	-10.5bp	-6.1bp
Panel B: Characteristics Linked to the ETF Ecosystem					
% ETF Ownership	5.2%	0.8%	4.6%	5.1%	5.7%
% Low ETF	15.5%	9.5%	8.3%	14.5%	21.4%
PT LCS	0.84	0.10	0.77	0.83	0.90

Panel C: Low ETF Ownership and Trade Activity

	(1) Trade Count	(2) Volume
Constant	92.4*** (10.70)	1.18e8*** (15.78)
Low ETF	-38.3*** (-3.38)	-3.02e7*** (-3.06)
LCS	118.3*** (12.44)	7.19e7*** (8.69)
Bond Observations	9,515	
Sample Period	January 1st 2021 – December 31st 2021	

Table XII: Relationship to the ETF Ecosystem – HY Portfolio-level Regressions

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

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: *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of IG or HY ETFs), *% low ETF* (the proportion of the portfolio that is in the bottom quartile of ETF ownership), and *PT LCS* (the weighted average of the LCS of the line-items in the portfolio). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Alt text: The table shows that the HY PT benefit increases with the proportion of the portfolio that is owned by ETFs and as the liquidity of the underlying line items declines.

	$EHS_p - \widehat{EHS}_p$	
	(1)	(2)
Const	1.00*** (-5.70)	-6.58*** (-5.70)
% ETF Ownership	-1.16*** (-6.07)	-
% Low ETF	-	0.074*** (4.68)
PT LCS	-6.26*** (-4.28)	-5.55*** (-3.76)
Sample	2,054 PTs	

Table XIII: 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 ETF* (past 30-day correlation between portfolio and LQD returns), *% Rightway ETF C/R* (% of line items that were rightway for the aggregate create or redeem basket of LQD on the execution date of the portfolio trade), *% ETF Ownership* (average % of the bonds' amount outstanding held by the broader universe of IG or HY ETFs), *% Low ETF* (% of line items that are in the bottom quartile of ETF ownership), *PT LCS* (the weighted average LCS of the line-items in the portfolio, measured in percent), *L-S Dummy* (dummy equal to one if the portfolio has both a long and a short leg) and *L-S Dummy Balanced* (dummy equal to one if the long and the short leg are balanced). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively. The continuous dependent variables are winsorized at the 5% level.

Alt text: The table shows that long-short PTs get similar or even better execution than uni-directional PTs, in contrast to the prediction of the diversification hypothesis.

	$EHS_p - \widehat{EHS}_p$					
	IG		HY			
	(1) All	(2) L-S only	(3) All	(4) L-S only	(5) All	(6) L-S only
Const	-3.44*** (-5.21)	-1.5** (-1.88)	0.61 (0.33)	9.41*** (2.97)	-6.92*** (-5.86)	-5.51*** (-3.19)
Corr ETF	-1.35* (-1.79)	-2.56*** (-2.78)	-	-	-	-
% Rightway ETF C/R	-5.32*** (-5.77)	-4.01*** (-2.83)	-	-	-	-
% ETF Ownership	-	-	-1.13*** (-5.91)	-2.05*** (-6.10)	-	-
% Low ETF	-	-	-	-	6.99*** (4.41)	9.18*** (3.86)
PT LCS	-0.04 (-0.08)	-3.16*** (-3.68)	-6.17*** (-4.21)	-10.41*** (-4.54)	-5.40*** (-3.66)	-6.97*** (-3.14)
L-S Dummy	-0.78*** (-2.89)	-	0.51* (1.80)	-	(0.47) (1.64)	-
L-S Dummy Balanced		-0.29 (-0.68)	-	-0.21 (-0.40)	-	0.01 (0.01)
Sample	1,987 PTs	808 L-S PTs	2,054 PTs	730 L-S PTs	2,054 PTs	730 L-S PTs

Table XIV: Alternative Explanations – Signalling

$$\boldsymbol{EHS_{i,j,t} = \beta_1 PT_{i,j,t} + \beta_2 Pressure_{i,j,t} + \beta_3 PT_{i,j,t} \times Pressure_{i,j,t} + \Gamma Z_{i,j,t} + \lambda_{j,t} + \epsilon_{i,j,t}}$$

The table reports transaction-level regressions of effective half spread on a portfolio trade dummy ($PT_{i,j,t}$) and an interaction term with the dummy $Pressure_{i,j,t}$, which equals to one if this trade and the last five trades were all buy trades or this trade and last five trades were all sell trades. 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_{i,j,t}$), previous trade *EHS* ($Lag\ EHS_{i,j,t}$), and a dummy equal to one if the trade is executed before 13:00 ES. All continuous variables are winsorized at the 1% level. All regressions include a full set of bond-date fixed effects ($\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.

Alt text: The table shows that bonds with high potential for adverse selection do not have a lower PT benefit.

	IG	HY
Portfolio Trade	-5.90*** (-24.05)	-10.56*** (-26.67)
Press Dummy	2.45*** (14.79)	7.09*** (18.81)
Portfolio Trade × Press Dummy	-3.52*** (-11.65)	-11.65*** (-15.91)
Block Trade	-2.92*** (-15.24)	-9.63*** (-43.60)
Lag EHS	-0.16*** (-66.78)	-0.03*** (-9.62)
Morning Trade	-0.42*** (-3.95)	-4.01*** (-18.20)
Bond-Date FE	YES	YES
Bond-trade Observations	4,868,961	1,650,802
Sample Period	Jan 1st 2021 – December 31st 2021	

Table XV: Alternative Explanations – Misallocation of PT Price

$$\boldsymbol{EHS_p - \widehat{EHS}_p = \alpha + \beta_k \%PT\ Only_{k,p} + \epsilon_p}$$

The table reports the results of cross-sectional regressions of actual (*EHS*) minus predicted ($\widehat{EHS}$) effective half spread of portfolio $\boldsymbol{p}$ on the notional weighted proportion of the portfolio that traded only in PTs on the day of execution (*%PT Only*). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

Alt text: The table shows that PT only bonds do not distort the PT benefit as measured using the baseline regression.

	IG	HY
Intercept	-5.05*** (-24.93)	-10.11*** (-32.57)
% PT Only	-0.016** (-2.00)	-0.01 (0.53)
Sample	4914 PTs	2054 PTs
Sample Period	Jan 1st 2021 – December 31st 2021	

Table XVI: Alternative Explanations – High vs. Low PT Days

The table reports summary statistics of daily PT volumes as a percentage of total TRACE volume (column (1)) and summary statistics for SST *EHS* on days with low PT volumes (below the median) (column (2)) and on days with high PT activity (above the median) (column (3)).

Alt text: The table shows that the distribution of the EHS of SSTs is very similar on days with high and low PT activity, in both IG and HY.

	(1) % of PT Volumes	SST EHS, bp	
		(2) Low PT Volumes	(3) High PT volumes
Panel A: IG			
Mean	8.5%	15.2bp	14.3bp
Std	3.9%	2.74bp	4.73bp
P25	5.7%	13.4bp	12.8bp
Median	7.8%	14.8bp	14.4bp
P75	10.9%	16.6bp	15.6bp
Panel B: HY			
Mean	8.5%	27.4bp	26.5bp
Std	4.9%	4.1bp	3.2bp
P25	5.1%	24.5bp	24.1bp
Median	7.6%	27.1bp	26.4bp
P75	10.8%	29.8bp	28.7bp
Panel A and Panel B are based on 250 trading days during the period Jan 1st 2021 – Dec 31st 2021			

Table XVII: IG Portfolio Strategies

The table shows a summary of the different strategies investors use when trading portfolios of bonds. We classify a portfolio as having a specific motivation if the portfolio Herfindahl score (HHI) along that dimension is at least 50% higher than the respective HHI of the Bloomberg IG Corporate Bond Index along the same dimension.

Alt text: The table shows that many PTs are concentrated in specific market segments, notably maturity and sector.

Type of Strategy	**% of IG PT Volume**
Maturity Motivation	35%
Sector Motivation	24%
Rating Motivation	13%

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*.