Trade Protocols, Transparency, and Liquidity – Theory and Evidence from the European Corporate Bond Market

Jeffrey Meli † Barclays, US

Zornitsa Todorova Barclays, UK

Andrea Diaz Barclays, UK

22nd May 2024

Abstract

We develop a model of market making with an endogenous choice of trade protocols, and demonstrate that with a high inventory cost dealers engage in both principal and agency trading. Transparency shifts more transactions into the (uncertain) agency protocol, and increases the bid-offer of principal trades with a sufficiently low probability of agency execution. We test these predictions with a novel database of European corporate bond transactions, exploiting two sources of exogenous variation in transparency. Transparency increases transaction costs for large trades and trades in older bonds, which are more difficult to "match", and vice-versa.

JEL Codes:

Keywords: transparency, liquidity, transaction costs, dealer inventory, ETFs

---

We thank Adam Kelleher, Vincent Grandjean, Oliver Haste and seminar participants at Barclays for helpful comments.

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

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

Conflict of Interest Disclosure Statement

Jeffrey Meli

I have nothing to disclose.

Zornitsa Todorova

I have nothing to disclose.

Andrea Diaz

I have nothing to disclose.

1. Introduction

In this article we develop and test a model of market making that incorporates important structural changes in the corporate bond market since the GFC. First, we endogenize the choice of trading protocol and link the emergence of agency trading to the post-crisis increase in dealer inventory costs. In our framework, dealers engage in principal trading when the cost of inventory is low but offer a menu of principal and agency trading when the cost of inventory is high. We uncover an inverse relationship between the transaction cost of principal trades and the probability executing an agency trade. Second, we demonstrate that the effect of transparency on liquidity depends on the market environment. When the cost of inventory is high and adverse selection is low, transparency increases the use of agency trading and raises the cost of principal trades for which agency trading is more difficult. We argue that this best describes the current market, as the rise of bond ETFs has reduced adverse selection. ETFs provide real-time information about the price of corporate credit risk that was lacking in the earlier era, when bond-level transactions were the only source of information about valuation. Finally, we empirically verify this prediction by exploiting two recent sources of exogenous variation in trade reporting in the European corporate bond market, which allow us to identify the causal effect of transparency on both transaction costs and the mix of principal and agency trading via a series of difference-in-difference regressions. Our results challenge the truism that transparency improves over-the-counter (OTC) liquidity, which is based on the pre-crisis studies.

Our key assumption, which differs from previous models, is that the bargaining power of a market maker vis-à-vis investors changes once it has a position in inventory. Similarly to prior studies (e.g., Back, Liu, & Teguia (2018)), we assume that dealers make take-it-or-leave-it bids to asset "sellers" who need liquidity. However, once the position is in inventory, we assume that potential "buyers" make take-it-or-leave-it bids to the dealer, based on knowledge of the dealer's cost of holding inventory, which allows them to share in the gains from trade.

When the cost of inventory is low, dealers capture most of the gains from trade. This leads to a high liquidity equilibrium, in which dealers provide immediacy only through principal trading. Agency trading is not viable because the dealer does not want to risk not finding a match; it prefers to hold any unsold bonds in inventory. Transparency increases liquidity, through either higher volumes or lower bid-offer. This comports with the conventional wisdom, which was established under the pre-GFC market conditions, and our model explains why agency trading was uncommon at that time.

A high inventory cost reduces the reservation price of the market maker, who can no longer afford to provide a high level of liquidity via principal trades. Instead, it reserves principal trading for motivated sellers, who are willing to accept lower prices from the dealer. This raises the prospect of agency trading, which allows the dealer to avoid the implications of its inventory cost and thus facilitate some transactions for less motivated sellers. We show that the dealer's optimal strategy is to offer the seller a menu designed to induce separation: immediacy at a low price (i.e., costly principal trading) or uncertain execution at a higher price (i.e., cheap agency trading). When the seller faces a large liquidity shock, it chooses immediacy, and vice-versa. However, the availability of agency trading forces the market maker to reduce the bid-offer of principal trades; it must pay above the seller's reservation price or the seller would always prefer the possibility of better execution via an agency trade. The linkage between agency trading and the cost of inventory, the use of agency trading as a means of separation, and the connection between the availability and effectiveness of agency trading and the cost of principal trading are all new insights. Our model explains the post-crisis increase in agency trading.

Low adverse selection creates an additional nuance; absent transparency, a buyer may overpay for some bonds, if the expected cost of doing so is below its share of the gains from trade. This raises the potential profits of the dealer, and thus allows it to provide greater liquidity, effectively undoing the effect of a high inventory cost. Transparency allows the buyer to accurately price bonds, and thus forces the dealer to internalize its full inventory cost. This alters the balance between trading protocols: with transparency some (certain) principal trades are replaced by (uncertain) agency trades. The principal trades that do occur come with elevated transaction costs, so long as the probability of executing an agency trade is sufficiently low, such that the premium paid by the market maker over the seller's reservation price is small. In other words, under the combination of high costs and low adverse selection, transparency drives more trades into the agency protocol, and the linkage between the option for agency trading on the one hand, and the price of immediacy via a principal trade on the other, results in differential effects of transparency depending on the probability of finding a match in the agency protocol.

This leads to our main hypothesis: under currently prevailing market conditions, the effect of transparency on liquidity will vary by transaction, depending on the likelihood of executing an agency trade. For transactions which are relatively easy to match, transparency will reduce bid-offer spreads. For those that are difficult to match, transparency will increase bid-offer spreads. We distinguish between these types of transactions via the size of the trade and the age of the bond. Larger trades and trades in older bonds are more difficult to match, and thus we expect the effect of transparency to be negative for those trades, but positive otherwise.

To test these predictions, we must compare recent trades with and without transparency. We turn to Europe, where the MiFiDII reforms that took effect in January, 2018 provide a unique setting to study the effects of transparency. MiFiDII introduced trade reporting for a wide set of asset classes, including for corporate bonds. While no consolidated tape exists, the trades are public, and we assemble a comprehensive set of dealer-to-client trades, accumulated from a large number of different voice and electronic venues. We use this dataset to estimate the bid-offer spread of round-trip transactions, distinguishing between principal and agency trades.

We exploit two sources of exogenous variation in the transparency of corporate bond trades executed in the EU and in the UK: Brexit, and a data issue that affected EU (but not UK) trade reporting for one quarter. [1] These allow us to isolate the effect of transparency by comparing trades in the same bonds but with different reporting obligations. The results align with our predictions. First, the effect of transparency on the bid-offer of principal trades depends on the size of the position and the age of the bond; for smaller trades and younger bonds, transparency reduces bid-offer (e.g., by 5% for small trades), but for larger trades and trades in older bonds, transparency increases bid-offer (e.g., by 15% for large trades). Second, we use the one-quarter disruption in EU reporting in 2022 to distinguish between "treated" bonds (those with a change in transparency) from "control" bonds (those with no change) in difference-in-difference regressions. We demonstrate that the proportion of agency trades in treated bonds declined (vis-à-vis control bonds) when transparency was disrupted, and then increased once transparency was restored.

Literature Review

We contribute to two strands of the literature. First, our work relates to theoretical models of transaction costs and liquidity in over-the-counter (OTC) markets, including Duffie, Gârleanu, & Pedersen, (2005), who consider the effect of search and bargaining on valuation in OTC markets, and Zhu (2012), who studies the effect of a repeat contact with the same buyer. A number of articles examine the effect of transparency on liquidity in search models, including Duffie, Dworczak, & Zhu (2017), and Vairo & Dworczak (2023).[2] The conclusions are mixed. Duffie et al. (2017) find that transparency (in the form of a published benchmark) typically increases liquidity, although the effect on market makers is ambiguous. Vairo & Dworczak, (2023) consider the effects of both post-trade and pre-trade transparency, and conclude that pre-trade transparency leads to more efficient outcomes than post-trade transparency.

None of these models includes multiple trading protocols, which is an increasingly common feature of many OTC markets, such as the corporate bond market (Goldstein &

---

[1] Note that an investor trades where domiciled; investors have no discretion as to the jurisdiction they are subject to.

[2] See also Glebkin, Yueshen, & Shen, (2022).

Hotchkiss, (2020), Choi, Huh, & Shin, (2023)). Some models assume that market makers have immediate access to an inter-dealer market, which obviates the need to explicitly model inventory. In others, they face no inventory or short constraints (e.g., Vairo & Dworczak, (2023)). Our work is closest to Back, Liu, & Teguia, (2018), who, similarly to us, abstract from search costs and model a single market maker. They assume the market maker faces an infinite cost of inventory, and so never holds any positions; despite this restriction they assume the market maker retains all the bargaining power with potential investors. [3] The authors find benefits of transparency.

We also contribute to the empirical literature on the relationship between liquidity and transparency in the corporate bond market., based on the introduction of the Trade Reporting and Compliance Engine (TRACE) in the US in 2002 (e.g. Edwards, Harris, & Piwowar, (2007); Bessembinder, Maxwell, & Venkataraman, (2006)); Goldstein, Hotchkiss, & Sirri, (2006); Bessembinder & Maxwell, (2008). The general conclusion of these articles is that transparency decreases transaction costs and increases trading volumes. Our paper uses data from November 2022 to September 2023 and provides fresh empirical evidence on how post-trade transparency impacts transaction costs under current market conditions.

2. Model

2.1. Motivation and key assumptions

The corporate bond market has changed significantly since the introduction of TRACE. One of the most important changes is the increased cost associated with dealer inventory. Post-crisis financial reforms (e.g., the Dodd-Frank Act and the Basel III framework) raised the capital charges associated with inventory and restricted market-

---

[3] Back et. al. motivate this assumption by proposing that an alternative available to the dealer is to run an auction for the asset, inducing Bertrand competition amongst the interested buyers. However, search costs can limit the ability to identify a sufficient number of buyers, and the costs of inventory are accrued over time. Our assumption is similar to that in models that assume some division of the gains from trade between the market maker and the buyer, such as Duffie et. al. (2005), although they do not explicitly include a market maker.

makers' risk-taking.[4] As a result, dealer inventory declined considerably after the crisis (***Figure 1***); as of September 2023, market-makers' corporate bond holdings were one twentieth their size in 2006, despite the market having grown significantly over that period.[5] A higher cost of inventory could conceivably alter the distribution of bargaining power between market makers and investors, and thus the provision of liquidity. For example, it seems intuitive that the increased post-crisis prominence of agency trading is linked to higher inventory costs. However, existing models of market making are not well-suited to explore this connection, nor how the use of agency trading interacts with principal trading.[6]

We assume a different dynamic between the buyer and the market maker, which allows for a richer interaction between inventory cost and liquidity and is based on conversations with bond traders.[7] In our model, the dealer makes a take-it-or-leave-it bid to a seller of an asset (i.e., the dealer has the bargaining power), but once the position is in its inventory, the buyer makes a take-it-or-leave-it bid to the dealer (i.e., the buyer has the bargaining power); this bid is conditioned on any information about the value of the security, the cost of inventory, and its effect on the dealer's reservation price. In other words, the cost of inventory determines the division of gains from trade between the dealer and the buyer.

Another important change is the rise of corporate bond ETFs, which have extremely high secondary market liquidity (Meli & Todorova (2023)) and provide real-time

---

[4] For example, Bao, O'Hara & Zhou, (2018) find that price impact increased for recently downgraded bonds more after the implementation of the Volcker rule compared to before; Dick-Nielsen & Rossi, (2018) use bond index restrictions as a quasi-natural experiment and find that the price of immediacy increased post-crisis versus pre-crisis; Adrian, Boyarchenko & Shachar, (2017) find that corporate bond liquidity provision declined significantly for market-makers that are more constrained by regulations.

[5] The figure is based on net positions in corporate bonds for US primary dealers, available through the Federal Reserve Bank of New York Primary Dealer Statistics database. Although similar data is not readily available in Europe, we expect a similar pattern exists for European market-makers.

[6] For example, Back, Liu, & Teguia, (2018) assume an infinite cost of inventory. Duffie, Dworczak, & Zhu, (2017)) assume a range of dealer costs, but that some non-zero measure of dealers have costs of zero (which they label "fast traders"). Neither is suited to assess the implications of different inventory costs, and in neither is there a choice of trading protocol.

[7] "We are more willing to negotiate with an investor who is taking us out of risk than when we are being asked to take a position into inventory", Head of IG Bond Trading at a large broker-dealer.

information about the price of corporate credit risk. When TRACE was introduced these instruments did not exist, and aggregated trade reporting was the only source of information about bond prices. The potential gains from transparency were large because potential investors were subject to a high degree of adverse selection. However, the number of ETFs, their AUM (***Figure 2***), and their secondary market liquidity has reduced the degree of adverse selection, potentially reducing the need for transparency, or at least muting its benefits for liquidity. We explicitly incorporate the degree of adverse selection into our model.

In our analysis we focus on the equilibria that obtain in a low inventory cost "pre-GFC" regime and a high cost and low adverse selection "post-GFC" regime.

2.2. Primitives

Players: There are three players, all of which are risk neutral and maximize expected payoff.

"Seller": An investor that owns a security of value v equal to either v_l or v_h with probabilities θ and $(1-\theta)$ respectively, and $v_l < v_h$. The seller experiences a liquidity shock Δ equal to Δ_u or Δ_d, with probabilities q and $(1-q)$ respectively, with $\Delta_u > \Delta_d$, such that the value of the security to the seller is $v - \Delta$; as is standard in these models, the liquidity shock is the source of gains from trade. The liquidity shock and the value of the security are uncorrelated.

"Dealer": A market maker in the security, willing to provide liquidity to the seller. The dealer can either sell the security once acquired (see below), or hold it in inventory, at a cost $c \geq 0$.

"Buyer": An investor who is potentially willing to buy the security at a price negotiated with the dealer. A buyer arrives with probability p.

We assume that the liquidity shock is known only to the seller. The value of the security is known to both the seller and the dealer. The information available to the buyer depends on the transparency regime. Absent transparency, the buyer knows only the distribution of values and the cost of inventory. However, when we impose transparency via trade reporting, the buyer can infer the value of the security from the transaction reporting. Finally, we assume that all players have a weak preference for trading.

Trading

Two trading protocols are available to the dealer. We label the first protocol “principal trading”; the dealer trades sequentially with the seller and the buyer, and holds any unsold positions in its inventory. The bargaining power of the dealer depends on whether it has the position in inventory. Before it acquires a position, the dealer has the bargaining power: it makes a take-it-or-leave-it offer to the seller. Once the dealer owns the security, the buyer has the bargaining power; if it arrives, the buyer purchases the security at the dealer’s reservation price, which is a function of the cost of inventory.

Second, we consider an “agency trading” protocol, in which the dealer attempts to pre-arrange both the sides of the trade and execute them simultaneously. Agency trading is only successful if a buyer arrives, which occurs with probability p. Therefore, agency trading requires the seller to sacrifice certainty. Because the dealer does not hold the security in inventory, and can refuse to execute the trade if its terms are not met, we assume that the dealer has bargaining power over the buyer in this protocol, meaning the security is sold at the buyer’s reservation price. Finally, we assume that the dealer cannot “lie” to the buyer about the terms of the transaction, which implies that in an agency trade the buyer learns the value of the security, even absent transparency.[8]

Dealers can utilize either protocol or provide a “menu” to sellers. As we will see below, a menu will involve a trade-off: certain, but expensive, principal trading, versus uncertain, but cheap, agency trading.

Timing and equilibrium definition

---

[8] We make this assumption for analytical convenience; our conclusions survive if the buyer does not learn the value of the security. In the context of our data, our assumption is more appropriate, because the European regime does eventually release trade details; reduced transparency only delays this revelation. Therefore, given the market standard that the dealer states its purchase price in an agency trade, it is unlikely that a dealer could sustainably prevaricate.

We use a one period model, with the following stages. First, the seller asks the dealer for a bid on the security, which is a function of the security value (because the dealer knows the value of the security). If the dealer purchases the security or accepts an order, it can sell the security to a buyer, if one arrives. We define the bid-ask spread to be the average of the difference between the purchase and sale prices for round-trip trades; this will vary by protocol. Finally, at the end of the period all players realize their payoffs, which for the dealer includes a cost incurred for any remaining inventory.

Equilibrium is defined as a set of strategies that constitute a Trembling Hand Perfect Equilibrium for each of the players.

2.3. Simple case: no buyer

We first consider a simple case where there is no buyer, meaning the dealer must hold any securities purchased in inventory. We drop the subscripts from v because the dealer knows the value of the security. Due to the discrete nature of the liquidity shock the dealer will make one of two bids: $v - \Delta_u$ or $v - \Delta_d$. In the former case, the trade is executed with probability q, i.e., when the seller experiences a large liquidity shock. In the latter case, the transaction is executed with probability 1. The dealer chooses the bid that maximizes its expected profit, equal to the value of the security less the sale price and the cost of inventory. In particular, the dealer will bid $v - \Delta_d$ when:

$$q * (\Delta_u - c) < \Delta_d - c \rightarrow$$

$$c < \frac{\Delta_d - q\Delta_u}{1-q} = c' \qquad [1]$$

We assume that $\Delta_d > q\Delta_u$. When the cost of inventory is low (or zero), the dealer bids $v - \Delta_d$. This is a "high liquidity" equilibrium, with high transaction volumes and prices that are close to "fair" value (we formalize bid-offer below when we include a buyer). The seller earns expected utility of $q(\Delta_u - \Delta_d)$ because it fills its liquidity need at a high price even when its liquidity shock is large. Dealer profits are $\Delta_d - c$.

When $c > c'$ we have a low liquidity equilibrium. Transaction volume declines to q, and the discount to fundamental value increases Δ_u. The seller earns utility of 0 because it does not trade when it has a low liquidity shock and trades at its reservation price when the liquidity shock is high. Dealer expected profit is $q(\Delta_u - c)$.

2.4. Adverse selection

Absent transparency, the buyer faces potential adverse selection, which will affect the price that it is willing to pay for a security in the dealer's inventory. Due to the discrete nature of security value, the buyer will make one of two bids for a security in dealer inventory: $v_h - c$ or $v_l - c$. The former offer is equal to the dealer's reservation price for the high value security, and exceeds the reservation price of the low value security. Therefore, the dealer sells both securities at that price. The dealer will only sell the low value security if the buyer makes the latter offer.

The buyer chooses the offer that maximizes its utility. While the higher offer entails overpaying for the low value security, it increases the probability of trade, and thus generates a benefit of c across more transactions. The lower offer is guaranteed to avoid overpaying, but results in fewer transactions. We classify the degree of adverse based on the optimal strategy for the buyer.

Definition: When the optimal strategy for a buyer bidding on a security in dealer inventory is $v_h - c$, *adverse selection is low; otherwise adverse selection is high.*

The size of the inventory cost and the levels and distribution of bond values all determine the degree of adverse selection. For example, using the unconditional distribution of securities in the population, the higher bid is optimal if:

$$\theta * c < (1-\theta) * c + \theta * (c - (v_h - v_l)) \rightarrow$$

$$c > \left[\frac{\theta}{1-\theta}\right] * (v_h - v_l) \qquad [2]$$

As we will see below, the actual distribution of securities in the dealer's inventory may differ from the unconditional population distribution, which affects the threshold for adverse selection.

2.5. Introducing a buyer and the choice of trading protocol

With a buyer, the dealer must choose what trading protocols to make available to the seller: principal trading only, a menu of principal and agency trading, or agency trading only. We illustrate the comparison of trading protocols in the transparent market (allowing us to once again drop the subscripts from v).

For any security in inventory, the buyer bids $v - c$, equal to the reservation price of the market maker. Therefore, the optimal principal-only strategy for the dealer depends on the level of the inventory cost; if $c < c'$, the optimal strategy is to bid $v - \Delta_d$, and if $c > c'$ the optimal strategy is to bid $v - \Delta_u$.

A menu of principal and agency trading takes the form of:

$$B(v) = v - K \textit{ with certainty or}$$

$$B(v) = v - X \textit{ if a match is found} \qquad [3]$$

The goal of the menu is to induce separation, whereby the seller chooses expensive principal trading when it faces a high liquidity shock, and vice-versa. This is optimal for the dealer if the increased profits from principal trading when the liquidity shock is high outweigh the decline in (expected) profits from agency trading when the liquidity shock is low. In order to induce separation, it must be the case that $X \leq \Delta_d$; otherwise the seller would not choose agency trading when it faced a low liquidity shock. However, this implies that the seller would earn positive expected utility from agency trading when it faces the large liquidity shock. Therefore, it cannot be the case that $K = \Delta_u$, as that leads to a utility of 0 for the seller when it faces a high liquidity shock (it is paid its reservation price), and the seller would prefer the positive expected utility from agency trading. In other words, the availability of agency trading necessarily reduces the cost of principal trading. Therefore, the dealer prefers X as large as possible, which reduces the required discount to Δ_u, implying that $X = \Delta_d$. Even so, it must reduce the cost of principal trading. Taking advantage of the weak preference for trading, the cost of principal trading must equalize the expected utility of the seller across the two protocols when the liquidity shock is large:

$$\varDelta_u - K = p(\varDelta_u - \varDelta_d) \rightarrow$$

$$K = (1 - p)\varDelta_u + p\varDelta_d \qquad [4]$$

The important insight in [4] is that the required discount to the cost of principal trading increases with the probability of a match in the agency market. For example, if $p = 1$, $K = \varDelta_d$, meaning that the dealer prices principal trades equivalently to agency trades.

Finally, we consider the agency-only strategy. In a successful agency trade, the dealer sells the security to the buyer at v (its reservation price), and thus earns profits equal to the difference between v and the price paid to the seller. Because we assume that $\varDelta_d > q\varDelta_u$, the dealer will pay the seller $v - \varDelta_d$, and this strategy earns expected profits of $p\varDelta_d$.

As we will see below, the degree of adverse selection (absent transparency) and the level of inventory costs will determine the optimal choice of protocol for the dealer.

2.6. Pre-GFC: low inventory cost

We first analyze the low cost pre-GFC equilibrium, and assume that $c < c'$. We start with the transparent market, using the insights from the prior section, and then examine how the equilibrium changes when we remove transparency.

Transparency

With a low inventory cost, the optimal principal-only trading strategy is to buy all bonds at a price of $v - \varDelta_d$, earning profits of $\varDelta_d - c$. This is preferable to the menu described in [3] when:

$$\varDelta_d - c > q * (K - c) + (1 - q) * \varDelta_d * p \quad \rightarrow$$

$$c < (1 - p)c' \qquad [5]$$

which we obtain by substituting for K from [4]. Further, if [5] is satisfied then the profits from principal trading are larger than the profits from agency trading ($p\varDelta_d$). Therefore, for sufficiently low inventory costs, the dealer provides liquidity only via principal trading, which comports with the observation that agency trading was relatively rare before the GFC.

Lemma 1: When inventory cost is low as in [5], the transparent equilibrium involves only principal trading.

This leads directly to the equilibrium in the low cost, transparent market (all proofs are in the Appendix).

Proposition 1: The unique low cost (as in [5]) transparent equilibrium is:

- *a) The dealer offers $v - \Delta_d$ to the seller, and buys all securities;*
- *b) The buyer (if it arrives) buys securities at $v - c$ from the dealer;*
- *c) Total transaction volume equals $1 + p$;*
- *d) Realized bid-offer on round trip trades equals $\Delta_d - c$.*

Trading volumes equal $1 + p$ because the dealer buys all securities on a principal basis and sells to the buyer, if it arrives. The transaction cost reflects the optimal bid of the market maker (i.e., as if the liquidity shock is low) and the inventory cost.

No Transparency

Absent transparency, agency trading remains unsustainable in equilibrium as long as [5] is satisfied. The profitability of a principal trade for the dealer is at least as high is the case with transparency, because the dealer is paid at least its reservation price on any security in inventory (the dealer can be paid above its reservation price on low value securities if adverse selection is low, increasingly the appeal of principal trading).

The resulting equilibria have one of two differences to the transparent equilibrium. When adverse selection is high, removing transparency reduces trading volumes, because trades in the high value securities are one-sided. Conversely, when adverse selection is low, removing transparency increases bid-offer, because the buyer overpays for low quality securities.

Proposition 2: The unique low cost (as in [5]) non-transparent equilibrium is:

a) *The dealer offers* $v - \Delta_d$ *to the seller, and buys all securities on a principle basis;*
b) *When adverse selection is high, the buyer (if it arrives) purchases low value securities at* $v_l - c$*, the dealer holds the high value securities in inventory, total transaction volume is* $1 + p * \theta$*, and realized bid-offer on round trip trades equals* $\Delta_d - c$*;*
c) *When adverse selection is low, the buyer (if it arrives) buys all securities at* $v_h - c$*, total transaction volume is* $1 + p$*, and realized bid-offer on round trip trades equals* $\Delta_d - c + \theta * (v_h - v_l)$*;*

We conclude that transparency increases liquidity when inventory costs are low, in keeping with the findings of the existing literature. It either raises volumes, when the degree of adverse selection is high enough to otherwise disincentivize some trading, or it reduces bid-offer, by keeping the buyer from overpaying for some securities. In the former case, transparency has no effect on the dealer, as the foregone transactions occur at its reservation price, nor on the seller, which sells all securities to the dealer regardless. However, the welfare of the buyer increases; it earns c on a larger number of trades. In the latter case, the seller is similarly indifferent, but the dealer profits decline with transparency, because it is not able to sell overpriced securities. Conversely, transparency benefits the buyer because it pays the dealer its reservation price on all transactions.

Corollary: When the cost of inventory is low transparency increases liquidity.

2.7. Post-GFC: High cost and low adverse selection

We now assume that the inventory cost is high ($c > c'$) and that adverse selection is low. We will demonstrate a high cost can change the implications of transparency; however, this will hold only in a specific range of costs above c'. We will identify that range and demonstrate how and why the effect of transparency differs in that range (outside that range transparency has the standard effect). As above, we start with the transparent market.

Transparency

We first compare the optimal high-cost principal-only strategy ($v - \mathit{\Delta}_u$, which trades with a probability q) with the menu of trading protocols outlined in [3] and [4]. This menu is preferred by the dealer when it generates greater profits:

$$q * (\mathit{\Delta}_u - c) < q * (K - c) + (1 - q) * \mathit{\Delta}_d * p \rightarrow$$

$$q\mathit{\Delta}_u < \mathit{\Delta}_d \qquad [6]$$

However, [6] is true by assumption (ensuring that low inventory costs lead to a high liquidity equilibrium). Therefore, the equilibrium cannot involve only principal trading. Similarly, an equilibrium with only agency trading is not viable if:

$$q * (K - c) + (1 - q) * \mathit{\Delta}_d * p > \mathit{\Delta}_d * p \rightarrow$$

$$(1 - p) * \mathit{\Delta}_u > \mathrm{c} \qquad [7]$$

Equation [7] implies that an extreme inventory cost renders all principal trading nonviable. However, so long as the inventory cost is not that high, the equilibrium will involve a menu of principal and agency trading. Therefore, we limit the cost to the range $\mathrm{c}' < c < (1 - p) * \mathit{\Delta}_u$, where the equilibrium involves a menu of principal and agency trading:

Proposition 3: The unique high cost, transparent equilibrium is:

- *a) The dealer offers the seller a choice of a certain principal trade at $v - K$ or an agency trade at $v - \mathit{\Delta}_d$ (which has success rate p), for K defined in [4];*
- *b) The seller chooses immediacy when it faces a large liquidity shock and the agency protocol when it faces a low liquidity shock;*
- *c) All securities in inventory are sold to the buyer at $v - c$;*
- *d) Total transaction volumes equal $q + p * (2 - q)$*
- *e) Realized bid-offer on round trip principal trades equals $K - c$.*

Volumes reflect the choice of protocol: the dealer purchases securities on a principal basis when the seller experiences a large liquidity shock and sells them to the buyer when it arrives. It also matches both sides of an agency trade with probability p when the seller experiences the small liquidity shock.

Agency trading plays a specific economic function: it facilitates trades that would otherwise be precluded by the high cost of inventory. However, its existence as an option has ramifications for principal trades. They must be done at a lower bid offer, in order to induce separation. The discount on principal trades grows as the probability of a match increases (and thus as K declines), because a higher probability of a match increases the seller's expected utility from the agency protocol when it has the large liquidity shock. Of course, the net effect of introducing agency trading is greater dealer profits; the dealer makes up for the lost revenue from lower-cost principal trades via the agency trades. The availability of agency trading increases the utility of the seller, who earns positive expected utility when it experiences the large liquidity shock.

No transparency

The assumption that adverse selection is low implies that the buyer is willing to pay the reservation price of the high value security. This increases the dealer's profit from the low value security, which can in turn affect the liquidity the dealer is willing to provide a seller of that security. In particular, if the dealer can sell the low value security at $v_h - c$, then the threshold computed in [1] no longer applies, because buying the low value security is so profitable. The new threshold is determined by:

$$q * [\Delta_u - c + (v_h - v_l)] < \Delta_d - c + (v_h - v_l) \rightarrow$$

$$c < \frac{\Delta_d - q\Delta_u}{1-q} + (v_h - v_l) = c' + (v_h - v_l) = c'' \qquad [8]$$

Note that the original threshold still applies to the high value security. This raises the intriguing possibility that when the cost of inventory is between the two thresholds (i.e., $c' < c < c''$) the dealer will provide differential liquidity for the two types of securities.

Of course, this would change the relative proportions of securities in the dealer inventory, which would contain all low value securities but only a portion of the high value securities (i.e., those where the seller experienced the large liquidity shock). In particular, the proportion of low value securities is equal to θ_l:

$$\theta_l = \frac{\theta}{[\theta + q * (1-\theta)]} > \theta \qquad [9]$$

To be an equilibrium, it must be the case that the buyer is still willing to purchase all securities at the high value even when accounting for the increased proportion of low value securities. In other words, adverse selection is worse for an inventory cost in this range, such that the constraint in [2] is actually tighter, because we substitute θ_l for θ:

$$c > \left[\frac{\theta_l}{1-\theta_l}\right] * (v_h - v_l) \qquad [10]$$

We now further restrict the range of inventory costs we consider:

$$\max\left(c', \left[\frac{\theta_l}{1-\theta_l}\right] * (v_h - v_l)\right) < c < \min\left(c'', (1-p) * \Delta_u\right) \qquad [11]$$

In this range, the dealer will behave as if its inventory cost is low when presented a low value security, and as if its inventory cost is high when it is presented a high value security. This implies that the dealer buys all low value securities on a principal basis, but offers the menu of agency and principal trading when presented the high value security. Further, the buyer purchases all securities in inventory at a price of $v_h - c$. In other words, the equilibrium is a hybrid, with the "low" liquidity equilibrium (i.e., that which includes both agency and principal trading) for the high value security and the "high" liquidity equilibrium (all trades executed on a principal basis) for the low value security:

Proposition 4: With cost between the two thresholds in [11] and no transparency, the unique equilibrium is:

a) *For the high value security the dealer offers the seller a choice of a certain principal at* $v_h - K$ *or an agency trade at* $v_h - \Delta_d$ *(which has success rate* p*), for K defined in [4];*
b) *For the high value security the seller chooses principal trading when it faces a large liquidity shock and agency trading when it faces a low liquidity shock;*
c) *The dealer buys all low value securities on a principal basis at* $v_l - \Delta_d$*;*
d) *The buyer (if it arrives) purchases all securities in dealer inventory at* $v_h - c$*;*
e) *Total transaction volume equals* $\theta(1+p) + (1-\theta) * (q - pq + 2p)]$

f) *Realized bid-offer on round trip principal trades equals* $\theta_l * [\Delta_d + (v_h - v_l)] + (1 - \theta_l) * K - c$.

Transaction volume reflects the fact that the dealer buys all low value securities and high value securities that are paired with a large liquidity shock on a principal basis, and sells these positions when a buyer arrives. Further, the dealer (sometimes) matches buyers and sellers of the high value security when the seller experiences the small liquidity shock. The average bid-offer on principal trades reflects the mix of securities bought through that protocol. A low value security starts with a low bid offer ($\Delta_d - c$) that is adjusted upwards due to adverse selection. A high value security starts with a high bid offer (reflecting the large liquidity shock) that is adjusted downwards to induce separation on the part of the seller.

We now analyse how transparency affects volumes, bid-offer, and the mix of trading protocols. Total volumes are higher without transparency when:

$$\theta(1+p) + (1-\theta) * (q - pq + 2p)] > q + p(2-q) \quad \rightarrow$$

$$1 > q + (1-q) * p \qquad [12]$$

This is always true (except when $q = 1$). Therefore, with high cost and low adverse selection, transparency reduces volumes. The reduction is linked to a change in the mix of protocols. Transparency causes dealers to replace some certain principal trades with uncertain agency trades; specifically, when a low value security is paired with a small liquidity shock, the trade is done on a principal basis absent transparency but on an agency basis with it. Transparency stops the dealer from pooling principal transactions at the high price, which is otherwise feasible when adverse selection is low. In other words, it reduces the bargaining power of the dealer vis-a-vis the buyer, which is reflected in the liquidity provided to the seller.

Second, we determine the circumstances in which transparency increases the average bid-offer on principal trades:

$$\theta_l * [\Delta_d + (v_h - v_l)] + (1 - \theta_l) * K - c < K - c \quad \rightarrow$$

$$p < 1 - (v_h - v_l)/(\Delta_u - \Delta_d) \quad [13]$$

According to [13], it is possible that the average bid-offer spread also increases with transparency. This requires that the probability of an agency match is sufficiently low. The probability of an agency match features in the average principal bid-offer because K declines as the probability of a match increases: the immediacy offered by a principal trade must come at a better price in order to induce separation when a match is likely. Both the changes in volumes and the increase in bid-offer for some types of transactions are driven by a shift in trading protocols: transparency increases the proportion of trades done in the agency trading protocol.

Corollary: When inventory cost is high, transparency can increase the average bid-offer of principal trades and raise the proportion of agency trades.

This requires both low adverse selection and a cost of inventory that is high enough to deter some principal trading, but not high enough to prevent the dealer from holding any inventory nor from taking advantage of low adverse selection to pool securities. Under these circumstances, it is still not the case that transparency always reduces liquidity. It is specifically trades that are more difficult to match via the agency protocol that suffer from transparency. Trades which are relatively easy to match benefit. This leads to our testable predictions:

H1: Transparency reduces bid-offer spreads for trades that are easy to match, but increases bid-offer spreads for trades are difficult to match.

H2: Transparency increases the proportion of agency trades.

3. Constructing a database of European corporate bond transactions

3.1. Transactions reporting in Europe

To test these predictions, we need a data set of corporate bond transactions with differential transparency in a modern setting. Transparency in the US market has not changed since TRACE was introduced over two decades ago, and so we turn to Europe, where the Markets in Financial Instruments Directive (MiFiDII) rules require that all transactions in corporate bonds executed in Europe are reported with a unique bond identifier, an exact execution timestamp, price and quantity.[9] As we will see below, several developments allow us to identify exogenous variation in reporting under the MiFiDII rules.

However, we still need to construct a comprehensive repository of dealer-to-customer corporate bond transactions. While technically this information is publicly available free of charge, a major practical limitation in using the data for research purposes is that a consolidated tape does not exist. Unlike in the US, where TRACE is a single centralized repository, European data are published across a large number of different reporting venues. This complicates the process of collecting, cleaning and aggregating transactions.

Transactions executed on online platforms (eg. Tradeweb or MarketAxess) are reported by the respective platform, while OTC voice transactions are disclosed through an Approved Publication Arrangement (APA), which acts as the reporting entity on behalf of market-makers.[10] Each market-maker has one unique APA, which publishes all of its voice transactions. Note that the majority of leading electronic platforms also operate a separate and independent APA – e.g., Tradeweb and Tradeweb APA; Bloomberg and Bloomberg APA. For simplicity, we refer to both electronic platforms and APAs as “venues”.

We first collect data from 14 trading venues (for a detailed list refer to ***Table A 1*** in the Appendix). We then aggregate and clean the data (e.g., remove duplicates, reversals and amendments, etc.). We focus on euro-denominated investment grade (IG) corporate bonds executed between November 2022 – September 2023.[11] For each transaction, we obtain the exact execution and reporting timestamp (which as we will see below are often

---

[9] For more details, refer to this report by ESMA.

[10] Data are made available on public websites, for examples see https://www.apa.tradeweb.com/; https://www.bloomberg.com/professional/product/apae/

[11] We collaborated with Propellant.digital to build this database.

different), the cash price, the size, the venue on which the transaction was executed, the Market Identifier Code (MIC) of the venue and the jurisdiction (EU or UK).

We supplement the transaction dataset with static data (issuer, sector, issue size) and time-varying bond attributes (remaining years to maturity, bond age, an average credit rating using S&P, Moody's and Fitch ratings, and amount outstanding), which we obtain from Bloomberg. This raw dataset spans 2.4 million transactions and a total of €2.2 trillion of volume. It contains more than 5,000 unique bonds and 1,000 unique issuers.

To evaluate the representativeness of the data, we also collected a second proprietary dataset of corporate bond request for quotes (RFQs) executed by the Barclays trading desk over the period November 2022 to May 2023. The database contains a mix of dealer-to-customer and dealer-to-dealer RFQs; however, for confidentiality reasons, the identity of the contra-party Barclays was facing is masked. Barclays is one of the largest market-makers with a significant presence in the fixed income space. Hence, it is reasonable to assume that the sample of Barclays RFQs is representative of the corporate bond market as a whole. A large overlap between the Barclays RFQs and the transaction dataset would indicate that the database we have constructed is representative of the European corporate bond market.

We were able to match between 85% and 90% (by count and by volume) of the Barclays RFQs to the transaction dataset. In conversations with the trading desk we have verified that the majority of the unmatched RFQs were executed on dealer-to-dealer electronic venues, which are not part of transaction dataset.[12] Further, while we don't have a precise estimate of the size of the wholesale corporate bond market in Europe, TRACE estimates[13] show that during the same time period, dealer-to-dealer activity in the US constituted c.15% of total volumes, which is closely aligned with our matching

---

[12] Leading venues in this category are TPICAP and BGC/GFI. Our dataset also does not capture Euronext and German exchanges. However, we do capture LSE.

[13] TRACE explicitly differentiates between dealer-to-client and dealer-to-dealer volumes. Nothing in the existing literature suggests a systematic difference between US and European dealer-to-dealer volumes.

rate. These tests give us confidence that the dataset we have constructed captures close to 100% of the dealer-to-client European corporate bond trades over the relevant period.

3.2. Measuring transaction costs – Imputed Roundtrip Cost (IRC)

We measure transaction costs using the imputed round-trip cost (IRC) (Feldhütter, (2012); Kargar, et al., (2021)).[14] To construct the IRC, we first identify pairs of round-trip trades. A round-trip trade consists of two matched trades in the same bond with the same trade size that are executed as close as possible to each other but have different prices.[15] When a trade has more than one match, we use the match closest in time. On an intuitive level, the goal of our methodology is to impute the direction of trades and, in so doing, identify a sale from an investor to a market-maker, and the subsequent purchase by another investor, or vice versa. For each round-trip trade, we calculate the *IRC* as the percentage difference between the higher and the lower price and report the values in basis points:

$$IRC = 10{,}000 \times \frac{(P_{max} - P_{min})}{P_{min}}$$

Higher (lower) values of *IRC* signify higher (lower) transaction costs, and hence lower (higher) liquidity. We remove zero-cost round-trips and *IRC* values above the 95th percentile of the distribution to ensure that our results are not polluted by extreme values.

Following the recent literature (Kargar, et al., (2021)), we explicitly differentiate between agency and principal round-trips. We identify agency round-trips as trades executed within 15 minutes, whereas principal round-trips are trades executed more than

---

[14] Another measure commonly used in the literature (Bessembinder (2003); Collin-Dufresne, Junge, & Trolle (2020); Hagströmer (2021)) is the effective half spread, which gives the distance between the traded price and a benchmark price (e.g., the mid-price), taking into account the direction of the trade (buy or sell). We cannot use the effective half spread because MiFiDII post-trade data does not report the direction of trades. Other transaction cost measures include e.g. Amihud's (2002) price impact or Roll's (1984) autocovariance in price returns, produce noisy estimates when applied at the transaction level.

[15] Goldstein & Hotchkiss (2020) refer to these types of round-trips as "paired round-trips". In other methodologies, trades are matched with more than trade in the same bond up to the face amount of the initial buy trade.

15 minutes apart.[16] Our final sample contains c.666K observations (roundtrips), of which c.630K are principal round-trips and c.36K are agency round-trips.

Robustness

We test robustness of the *IRC* methodology by comparing the transaction costs of agency and principal round-trips, and the transaction cost estimates produced by the *IRC* methodology to the Barclays Liquidity Cost Score (LCS).

Since market-makers do not use their balance sheet when they intermediate agency trades, they should cost less than principal trades. In our sample agency round-trips cost on average 17.2bp compared to 38.3bp for principal round-trips (***Figure 3***), which is closely aligned with the findings of Kargar, et al., (2021) for the US corporate bond market. We also find that within principal round-trips, transaction costs increase the longer it took a market-maker to find the other side of the trade. For example, round-trips where it took the market maker between 1-5 days to find the other side cost 30.1bp compared to 56bp for trades where the market maker closed the position after more than 10 days.

Second, we aggregate the *IRC* to the bond-month level and compare the estimates to LCS. LCS is a commercially available measure of transaction cost computed using quotes from the Barclays trading desk. It follows the methodology by Konstantinovsky, Yuen Ng, and Phelps (2016). LCS measures the transaction cost for an institutional-size trade, expressed as a percentage of the bond's price (hence higher LCS signifies lower liquidity). *IRC* closely tracks LCS (***Figure 4***).

4. Exogenous variation in transactions reporting

The trade reporting rules in MiFiDII are more complex than those in the US, and allow for both real-time and delayed reporting depending on bond and trade characteristics. In this section we outline the rules governing real-time versus delayed

---

[16] In robustness checks, we have used different thresholds (5minutes or 10 minutes) and have obtained very similar results.

reporting, and identify two sources of exogenous variation in reporting over our sample period that function as quasi-natural experiments for the effect of transparency on liquidity.

4.1. Transaction reporting rules

As a general matter, MiFiDII requires that transactions be reported as close as reasonably possible to real-time. However, the rules contain a series of exceptions which qualify certain transactions for a reporting delay of up to four weeks. The most important features that determine if a transaction qualifies for a delay are bond liquidity, trade size and inclusion in a package trade.

As a first step, the European Securities and Markets Authority (ESMA) makes a recommendation to the National Competent Authority (NCA) in each country regarding the trade characteristics that determine reporting (***Figure 5***).[17] ESMA makes a liquidity assessment for each bond and recommendations regarding the relevant size thresholds. Liquidity assessments are performed each quarter and the results apply to the next quarter. Every bond is classified as either “liquid” or “illiquid”, based on the recent history of trades in that bond.[18] Each year, ESMA also sets two global trade size thresholds.[19] Over the period that we study (Nov-2022 to Sept-2023), the thresholds were €2 million and at €3.5 million.

For liquid bonds, the reporting requirement depends on the size of the trade. If the trade size is below the two size thresholds, then the transaction must be reported in real-time; if the trade size is above either of the thresholds, reporting can be delayed up to four weeks. All trades in illiquid bonds can be reported with a delay of up to four weeks.

---

[17] For example, Authority for the Financial Markets (AFM) in the Netherlands or Bundesanstalt für Finanzdienstleistungsaufsicht (Federal Financial Supervisory Authority) in Germany. For a detailed list of the supervisory contact points in each country, refer to this document.

[18] It classifies a bond as liquid if it fulfils three conditions: 1) the daily traded notional is larger than €100K; 2) the daily average number of trades is greater than two; and 3) if it has been traded on at least 80% of the days in a given quarter. In practice, this definition applies only to recently issued bonds.

[19] These are the so-called "size specific to instrument" (SSTI) and "large in size" (LIS). SSTI and LIS are set at the 80th and at the 90th percentile of the trade size distribution.

Importantly, the ultimate determination of which trades qualify for a delay lies with the NCAs. Each NCA decides which of two size thresholds apply to trades in its jurisdiction, and can choose to override the bond liquidity classification recommended by ESMA or to extend further the reporting deferral. In practice, the reporting of virtually all transactions that qualify for a delay is in fact delayed for the full four weeks.

Finally, a transaction in a liquid bond can also be deferred if it was executed as a part of package trade (TPAC), where at least one of the instruments in the package is illiquid. Package trades are *"…composed of two or more instruments that are priced as a single unit, simultaneously executed, and where the execution of each component is contingent on the execution of all other components*".[20] Package trades are typically done for risk management and hedging purposes; for example, when an investor trades a corporate bond and a credit default swap at the same time.[21]

Our trade dataset includes both an execution timestamp and a reporting timestamp; we can identify which transactions were reported with and without a delay by comparing these. Further, when a transaction is delayed, the justification for the delay must be disclosed (column "Flag").

Our toy example in ***Table 1*** consists of four transactions in two unique bonds: *ABC* is liquid and *XYZ* is illiquid. The first transaction in bond *ABC* was reported without a delay; the second was delayed because it was a large transaction, whereas the third was delayed because it was part of a package trade (TPAC flag), despite the fact that it was a small transaction in a liquid bond. Bond *XYZ* was illiquid (ILQD flag), so all transactions in that bond would typically be delayed.

---

[20] Refer to ESMA's guidelines on the treatment of TPACs.

[21] Although the formal definitions are somewhat similar, a package trade is not equivalent to a portfolio trade. A package trade involves instruments from several asset classes, where a portfolio trade contains only corporate bonds.

4.2. Exogenous variation N.1 - Brexit

Before Brexit, the sole responsibility to perform liquidity assessments and to make recommendations for transaction deferrals lay with ESMA. This meant that each quarter each bond had a unique liquidity classification (liquid or illiquid) and unique thresholds separating small from large transactions. All transactions had the same reporting schedule, irrespective of whether they were executed in the European Union (EU) or in the UK.

After Brexit, the authority to delay reporting for transactions executed on UK venues was transferred to the Financial Conduct Authority (FCA), while ESMA retained its remit over transactions executed in Europe. While ESMA and FCA continued to follow the same process and use the same rules, their calculations are based on data collected from the trading venues under their respective jurisdictions. This generated two sources of exogenous variation at the bond and at the transaction level. First, the same bond could have two different liquidity classifications during the same quarter – it can be liquid according to ESMA and not eligible for a reporting deferral, and illiquid according to FCA and eligible for a deferral, or vice versa. Second, the same bond could have different size thresholds in the EU and in the UK, implying that the same transaction could be eligible for a deferral based on size in the UK but not in the EU.[22]

To demonstrate, in ***Figure 6*** we plot the percentage of transactions reported in real-time for bond-quarters classified as liquid by ESMA. We bucket transactions based on trade size and show the respective number for each trade bucket. We would expect transactions in liquid bonds below the size thresholds to be reported real-time since the reporting cannot be deferred. However, within each size bucket, we find that a substantial percentage of transactions are in fact reported with a delay. For instance, 28% of the

---

[22] Under some circumstances it is possible to override the SSTI/LIS and set the so-called "threshold floor" of €200,000 to a subset of bonds. Typically, this happens if a regulator deems that they don't have sufficient information for a given bond to assess whether the proposed global size parameters are appropriate. Hence, the same bond could have different size thresholds in the two jurisdictions – in other words, it could have the standard SSTI and LIS size thresholds in the EU, but the threshold floor in the UK. In that case, for instance, we can find a €300K transaction in a liquid bond reported in real-time by an EU venue while it is reported with a delay by a UK venue.

transactions smaller than €500K are reported with a delay (dark blue bars in ***Figure 6***). The jurisdiction effect (i.e., a different liquidity classification and/or a different size threshold in the EU and in the UK, as indicated by the green bars) accounts for the majority of those delays. Another (albeit small) portion of the variation can be explained by package transactions (TPAC). The remaining small percentage of the variation can be attributed to reporting errors or differences in requirements at the NCA level.

Investor rules

Post-Brexit rules not only impacted the reporting schedule of corporate bonds, but they also put restrictions on which legal entities investors were allowed to trade with. Before Brexit, most leading trading venues (e.g. Tradeweb, MarketAxess etc.) served all of their European clients through a single entity, typically domiciled in the UK. For example, Tradeweb operated through Tradeweb Europe Limited – a London-based investment firm, regulated by the FCA. Post-Brexit, trading venues were required to stand up independent and fully functional entities regulated within the EU. For example, in 2017 Tradeweb established Tradeweb EU BV and MarketAxess established MarketAxess NL B.V., both of which are based in Amsterdam and are regulated by the Dutch National Competent Authority.

As a consequence, post-Brexit, investors must now face the trading venue domiciled in their jurisdiction. For example, in order to be eligible to trade with the UK entity of Tradeweb, an investor must be "*…authorised in the United Kingdom as an investment firm, a credit institution or as a UK branch of a non-UK investment firm or credit institution…*".[23] Similarly, in order to be eligible to trade with the EU entity of Tradeweb, an investor must be "*…authorised under MiFID II, a credit institution authorised under EU Directive 2013/36/EU or an EU branch of a non-EU investment firm or credit institution…*".[24] This means that for the same transaction in the same bond, a UK investor is required by law to trade with a UK venue, whereas a EU investor must trade with a EU

---

[23] Tradeweb UK's Rulebook.

[24] Tradeweb EU's Rulebook.

venue. Note that nothing in these rules prevents investors from shopping for "best execution" across the list of venues which are legally allowed to operate in their jurisdiction (e.g. Tradeweb UK vs. MarketAxess UK). However, a UK investor cannot choose to trade with a EU entity and vice versa. In other words, the variation in reporting delays driven by the jurisdiction effect are exogenous; investors cannot determine the reporting, it is imposed on them based on their location.

4.3. Exogenous variation N.2 – temporarily reduced EU transparency

On the 19th October 2022, ESMA announced that it will not publish the next-quarter bond liquidity assessment due to a data quality issue.[25] In accordance with the MiFiDII playbook, all bonds for which no liquidity assessment had been published were deemed illiquid[26] from 16th November 2022 until the application of the next liquidity assessment on the 16th February 2023. Therefore, all transactions in these illiquid bonds automatically qualified for a reporting delay. ESMA was explicit in its press release that the only exception was newly issued bonds, which maintained their liquid status and did not qualify for a delay.

For similar reasons, the FCA also did not publish a liquidity assessment for the period from 16th November 2022 until the 16th March 2023. However, differently to ESMA, the FCA did not make a formal press release. As a result, the two jurisdictions responded in very different ways to the "no publication" event.

In the EU, the number of transactions reported with transparency decreased sharply on the 16th November 2022 and subsequently recovered on the 16th February 2023, when the next regular publication period began and the new reporting rules applied (Panel A, ***Figure*** *7*). In the UK on the other hand, the number of transactions reported with transparency was unchanged before and after the "no publication" event.

Further, the average age of bonds that were reported with transparency in the EU dropped from 2.5 years to 0.5 years precisely on the 16th November 2023, which is

---

[25] The press release can be found here.

[26] This is in line with Q&A 10 of section 4 of the MiFID II transparency Q&As.

consistent with ESMA's guidance regarding newly issued bonds (Panel B, ***Figure*** *7*). Again, there was no corresponding effect in the UK. Our analysis shows that the UK venues most likely applied the last published classification (i.e., the classification used for the period 16th August 2022 to 16th November 2022) for bonds issued before November 16th and reported all bonds issued during the "no publication" period with transparency.

This "no publication" event generated two additional sources of variation:

1. **Across jurisdictions**: the reporting of some bonds changed in the EU, but it did not change in the UK. This applied to bonds aged between 6 months and 2 years – i.e., the difference between the light blue and dark blue lines in Panel B of ***Figure*** *7*.
2. **Within the EU**: the reporting of some of the bonds traded in the EU changed ("treated" bonds), whereas reporting remained unchanged for others ("control" bonds). Due to the unique setting and the timing of this quasi-natural experiment, we are able to study the effect of transparency on bond liquidity twice –as treated bonds both enter and exit the "no publication" period.

4.4. Empirical Design

We use this exogenous variation in reporting in two ways. In ***Section 5***, we pool all the round-trip trades, and exploit variation in reporting at the transaction-level in the EU and in the UK. We focus on similarly sized trades where reporting varied due to a combination of Brexit effects during quarters with regular liquidity publications and the differential response of the EU and the UK during the "no publication" quarter. In ***Section 6***, we use variation at the bond-level in the EU generated by the differential treatment of bonds depending on their age in the EU during the "no publication" quarter. These latter tests are traditional difference-in-difference specifications.

5. Transparency and the probability of an agency "match"

5.1. Summary statistics

We compare the EU and the UK corporate bond market along several key dimensions and present the results in ***Table 2***. The EU is a bigger market, both in terms of number of transactions and total volume. Importantly for our analysis, investors trade the same bonds and in similar trade sizes in both jurisdictions; 98% of the bonds in our sample trade in both markets, and the distribution of bond characteristics is very similar. Further, transactions costs across the two markets are very comparable. On average, EU and UK investors pay the same *IRC* to trade the same bond (***Table 3***). Nonetheless, to address any selection bias, we exclude the small number of bonds which are never reported in real-time and which only trade in one jurisdiction. Hence, any difference in *IRC* we find for trades reported with and without transparency must be due to differences in the transparency regime.

5.2. Econometric model

We compare the transaction costs of round-trip *i* in bond *j* executed in jurisdiction *k* on day *t* when a round-trip is initiated with and without transparency in a formal regression model at the transaction level:

$$IRC_{i,j,k,t} = \beta_1 Transparency_{i,j,k,t} + \Gamma X_{i,j,k,t} + \Phi Z_{j,t} + \lambda_j + \delta_k + \gamma_t + \epsilon_{i,j,t} \quad (\textbf{\textit{Model 1}})$$

where $Transparency_{i,j,k,t}$ is a dummy variable equal to one if the first leg of round-trip *i* is reported with transparency (i.e., without a delay). The main coefficient of interest in ***Model 1*** is β_1, which gives the difference between the transaction-cost of roundtrips reported with and without transparency. If transparency reduces transaction costs, we expect $\beta_1 < 0$. The identification of the estimates comes from variation in the transaction costs of bonds which have a different reporting schedule in different jurisdictions.

We include round-trip level controls collected in the vector $\mathrm{X}_{i,j,\mathrm{k},t}$. These include: the number of days it takes to close a position, an electronic trade dummy and a package trade dummy. ***Figure 4*** shows that *IRC* increases the longer it takes a market-maker to close a position, which could bias β_1 if transparency also affects the inventory holding period. Anecdotal evidence suggests that electronic venues have better reporting discipline and commit fewer reporting errors than APAs, which report voice transactions.

We also control for time-varying bond characteristics ($Z_{j,t}$) such as bond age (years since issuance), the logarithm of amount outstanding, remaining years to maturity and a credit rating dummy. The purpose of including these controls is to isolate the effect of transparency on transaction costs from the effect of other bond characteristics which independently drive transaction costs. Finally, we also include bond (λ_j), jurisdiction (δ_k) and date (γ_t) fixed effects (for the first leg of the round-trip) to account for any (potentially unobservable) factors that could affect our results. We estimate the model for principal trades, and use agency trades in robustness checks.

5.3. The effect of transparency on transaction costs

Over the full sample, we find that transparency reduces transaction costs for principal trades (column (1) in ***Table 4***). All else equal, the average transaction cost of a principal trade reported with transparency is 1.4bp cheaper than the same trade when reported with a delay. Given an average *IRC* for principal round-trips of 38.3bp, the effect translates into a 3.7% reduction in transaction costs.

Our theoretical model predicts that transparency could increase transaction costs for trades that are more difficult to match. We proxy difficult-to-match trades in two ways: by size and by age. Corporate bonds trade infrequently, and typically have low turnover, which is why it is substantially less difficult for a market-maker to offload a €500K position compared to a €2M position. Further, bonds trade very frequently shortly after they are issued, after which their liquidity sharply declines. For example, average monthly turnover decreases from 25% for newly issued bonds to less than 5% for bonds issued more than five years ago. (***Figure A 1*** in the Appendix). Hence, recycling a large position and/or a position in an aged bond is more difficult.

To examine the differential effect of transparency across trades we augment ***Model 1*** by including the following interaction terms:

$$IRC_{i,j,k,t} = \beta_1 Transparency_{i,j,k,t} + \beta_2 Transparency_{i,j,k,t} \times Difficult\ to\ Match_{i,j,k,t} + \beta_3 Difficult\ to\ Match_{i,j,k,t} + \Gamma X_{i,j,k,t} + \Phi Z_{j,t} + \lambda_j + \delta_k + \gamma_t + \epsilon_{i,j,t} \quad \textbf{(Model 2)}$$

where $Difficult\ to\ Match_{i,j,k,t}$ is either trade size buckets or bond age. In the former case, we define four buckets: ≤€500K, (€500K-€1M], (€1M-€2M] and (€2M-M3.5M], where the ≤€500K category is our reference bucket; hence the effect of transparency for trades smaller than €500K is given by β_1 and the effect for trades larger than €500K is given by the sum of β_1 and β_2. In the latter, we use a continuous measure of bond age.

The effect of transparency for principal trades varies with trade size in the way we expect. Transparency decreases transaction costs for small trades and increases transaction costs for large trades (as evidenced by the statically significant and positive β_2 in column (2), ***Table 4***). In ***Figure 8*** we show the total effect of transparency on transaction costs by size bucket. Transparency increases transaction costs for trades in the (€2M-€3.5M] size bucket by 8.9bp, which translates to a c.15% increase. One feature of the corporate bond market is that the number of trades and the value of trades are not uniformly distributed by trade size. For example, trade sizes smaller than €500K account for 80% of the observations but only 20% of the total notional traded. Conversely, large trades account for a small number of the total number of observations but generate most of the volume (***Figure A 2*** in the Appendix). Weighing the effect of transparency for each trade size bucket shown on ***Figure 8*** by its contribution to total volumes, we calculate that on average, the effect of transparency is a 6% increase in transaction cost for principal trade (in contrast to the unweighted estimates in **Model 1**).

We obtain directionally similar conclusions when we look at the interaction between transparency and bond age (column (3), ***Table 4***). To demonstrate that these effects apply independently of those of trade size, we also retain the interaction terms with trade size buckets in the regression specification. The benefits of transparency apply less to older bonds ($\beta_2 > 0$). For example, transparency decreases transaction costs for newly issued bonds by 6%, but only by 2% for bonds issued between five and ten years ago (***Figure 9***). Transparency increases bid-offer for bonds that are more than ten years old (although there are very few transactions in that age bucket).

5.4. Robustness

A large portion of our round-trips are not completed on the same day. Therefore, it is possible that bid-offer spreads capture changes in market conditions which occur while the market maker is looking for the other side of the trade (e.g., if the market maker buys a large trade in a bond reported with transparency on a day when volatility is low, but offloads the position on a day when volatility is high). We ensure that the differences in bid-offer that we measure are related to transparency rather than changes in market conditions by including two-way date-fixed effects (i.e., a date dummy for the first and second leg of the round-trip) in column (1) of ***Table 5***; our results remain unchanged.

We also verify that differences in trading venue are not responsible for our results (e.g., if for whatever reason some venues have both higher transaction costs and are more likely to report trades with a delay) by including venue fixed effects in column (2) of ***Table 5***. We find very similar results to the those reported in ***Table 4***. Similarly, the results are unchanged if we include jurisdiction-date fixed effects, which control for market events affecting a specific jurisdiction on a given day (column (3)). We also obtain similar results when we limit our sample to round-trips where both legs are in the same jurisdiction (column (4), ***Table 5***).

Finally, we estimate the same model for agency trades and find that transparency has no impact on bid-offer spreads, which is aligned with the intuition of our theoretical model (columns (5) and (6), ***Table 5***). Agency trades are pre-negotiated, meaning that at the time the first leg of the round-trip was reported, the market maker had found a matching buyer. Hence, transparency should have no implications for transaction costs.

6. “No publication” quasi-natural experiment

6.1. Treated and controls

The “no publication” event caused an exogenous change in the reporting schedule of a subset of the bonds in our sample, whereas reporting remained unchanged for others. Due to the unique setting and the timing of this quasi-natural experiment, we are able to study the effect of transparency on bond liquidity twice, as bonds both enter and exit the "no publication" period:

- **Entering the "no-publication" period.** We define control bonds as those issued at most three months before the 16th November.[27] These bonds were reported with transparency both before and after that date. We define treated bonds as those issued between three and six months before 16th November. These were liquid enough to be reported with transparency before that date, but old enough to be reported with a delay afterwards.
- **Exiting the "no-publication" period.** We define control bonds as those issued at most six months before 16th February and which remain classified as liquid afterwards. These were reported with transparency both before and after that date. We define treated bonds as those issued between six months and three years before 16th February and classified as liquid afterwards. These were reported with a delay before 16th February and without a delay afterwards.

In ***Table A 2*** we verify that transactions in treated and control bonds were in fact reported as expected. 89.5% of transactions in treated bonds were reported with transparency before the 16th November 2022, and none were reported with transparency afterwards. Similarly, no transactions in treated bonds were reported with transparency before the 16th February 2023 and 98.4% were reported with transparency afterwards. Close to 90% of transactions in control bonds were reported with transparency in all cases.

6.2. Agency versus principal trading

We compute the proportion of agency trades for treated and control bonds shortly before and after they enter and exit the "no publication" period (***Figure 10***). Control bonds do not change reporting status, as they are always reported with transparency; as expected, the proportion of agency trading for these bonds remains unchanged both as they enter and exit the quasi-natural experiment window (***Figure A 3*** in the Appendix

---

[27] We exclude bonds issued in the last month, as these bonds have very different liquidity (both volumes and bid-offer) compared to bonds that have aged for a couple of weeks.

verifies the parallel trends assumption with daily data). However, as treated bonds enter the "no publication" event (and thus are no longer reported with transparency) the proportion of agency trading drops by roughly half, from 10% to 5.9%. We obtain the mirror image on the other side of the event window. As treated bonds exit the "no publication" event (and thus are once again reported with transparency), their proportion of agency trading increases from 8% to 15.4%. These results align with the predictions of our theoretical model.

6.3. Difference-in-differences transaction cost estimates

We remedy any specification concerns regarding our earlier results using a difference-in differences (DID) regression applied to the "no reporting" period. The DID approach compares a treated bond before and after the "no publication" event and compares its bid-offer to that of a similar control bond. The outcome of the control bond provides the counterfactual scenario; in other words, this is how the treated bond would have behaved in absence of the treatment.

We compare the transactions costs of roundtrip (*i*) executed in the EU in treated and control bonds (*j*), before and after the event date (*t*) in the following DID specification:

$$IRC_{i,j,t} = \alpha + \beta_1 Treated_{i,j} \times Post_t + \beta_2 Treated_j + \beta_3 Post_t + \Gamma X_{i,j,t} + \Phi Z_{j,t} + \epsilon_{i,j,t}$$

(**Model 3**)

The coefficient β_1 is the DID estimate, which gives the difference in IRC between treated and control bonds, before and after the event start date. We use the same transaction-level ($X_{i,j,t}$) and bond-date level ($Z_{j,t}$) controls as in our baseline transparency regression. We estimate two sets of conceptually equivalent DID regressions, as bonds enter and exit the "no publication" window (***Table 8***).

The DID estimate relies on two assumptions. First, that treated and controls are similar, and second that treated and controls are on parallel trends prior to the treatment. In both specifications β_2 (i.e., the difference between treated and controls) is economically small and/or statistically insignificant. Conditional on observable characteristics, treated and controls are similar, which we would expect given that only a slight difference in age separates the two categories of bonds. ***Figure A 4*** in the Appendix verifies that the parallel trends assumption holds.

As treated bonds switch from transparency to no transparency on the 16th November 2022, their bid-offer on principal trades decreases by 6 bp ($\beta_1 < 0$). Similarly, as treated bonds switch from no transparency to transparency on the 16th February 2023, their bid-offer on principal trades increases by 3 bp ($\beta_1 > 0$). Both results are statistically significant support our earlier conclusion that transparency can be costly.

We can also augment **Model 3** with a triple-interaction term designed to measure the ease with which a match can be found for an agency trade:

$$IRC_{i,j,t,s} = \alpha + \beta_1 Treated_{i,j} \times Post_t + \beta_2 Treated_{i,j} \times Post_t \times Match_s + \beta_3 Treated_{i,j} \times Match_s + \beta_4 Treated_j + \beta_5 Post_t + \Gamma X_{i,j,t,s} + \Phi Z_{j,t} + \epsilon_{i,j,t,s} \text{ (\textbf{Model 4})}$$

We use the (€1M-€2M] trade size bucket from above and credit rating (defined as being rated BBB) to proxy for the ease with which a bond can be matched in the agency protocol. In all specifications the signs and magnitudes of β_1 and β_2 match qualitatively our baseline results. Transparency is more costly for larger trades and for trades in lower rated bonds. However, the triple interaction coefficient is not statistically significant for the trade size variable, which is likely due to the smaller statistical power of these tests. We have identified 69 treated and 150 control bonds, which combined with the fact that bonds don't trade frequently limits the sample available for inference.

7. Discussion and policy implications

7.1. Forthcoming Changes to the EU and UK Trade Reporting

In an effort to enhance market data transparency and reduce fragmentation, regulators in both the EU and in the UK have recently published proposals to amend the existing framework for reporting corporate bond transactions.[28] Although the specific provisions

---

[28] In June 2023 representatives of the European Commission, the European Council and the European Parliament reached a political agreement on the MiFiDII/MiFiR review. Legislative changes are expected

and technical details differ slightly, the overarching goal in both jurisdictions is to implement faster disclosures for corporate bond trades and establish a consolidated tape, which will provide a single reference source of information for prices and volume of traded bonds. Our evaluation of these proposals suggests that the new rules will significantly increase transparency in Europe.[29] We estimate that the number of transactions reported in real-time will increase from 8% to c.80%.

In drafting these proposals, regulators have cited the existing literature based on TRACE data, which concludes that transparency unequivocally improves liquidity for all corporate bonds.[30] One of the contributions of our paper is to show that, under the currently prevailing market conditions, the effects of transparency are heterogeneous and jointly depend on a combination of trade and bond characteristics.

While proposals to overhaul the current reporting system have been set in motion in both the EU and in the UK, the exact details are yet to be disclosed. Our work supports a sliding transparency design with different reporting categories and different deferral periods, depending on the characteristics of the bonds and the trades. There exist multiple ways in which the same level of transparency can be achieved, but the market implications of exactly which types of trades are made transparent might vary widely. Policy makers and regulators are faced with a difficult optimization problem – maximize transparency subject to preserving liquidity. Our results suggest that small trades and trades in newly-issued bonds benefit from transparency, but the new framework should make provisions to protect larger trades, trades in older bonds or trades in lower-rated bonds. While the optimal design of the new framework is beyond the scope of this paper, it is a promising avenue for future research.

---

to come into effect in 2024. In December 2023, the UK's financial regulator, the FCA, published a consultation paper inviting market participants for comments and suggestions on a proposal to improve the transparency regime in the UK.

[29] For more details on changes to the EU rules, refer to this draft report; for details on changes to the UK rules, refer to Chapter 6 of the FCA consultation paper.

[30] Edwards, Harris, & Piwowar (2007); Bessembinder, Maxwell, & Venkataraman (2006); Goldstein, Hotchkiss, & Sirri (2006)).

7.2. Welfare implications

Our results do not imply that introducing transparency will have a net-negative effect on the European corporate bond market. Both our theoretical model and our empirical test suggest that welfare implications will depend on bond and trade characteristics, on the type of investor, and possibly on market conditions. Transparency benefits the smallest trade sizes the most and, by extension, retail investors who are more likely to trade these smaller tickets and who typically don't have access to timely, high-quality pricing data.

On the other hand, transparency decreases liquidity for the largest and most difficult to match trades. Institutional investors, which are more likely to trade larger tickets, might face higher transaction costs. This is a particular concern when they are forced to sell quickly and require immediacy (e.g., when they face a sizeable outflow, which we model as a large liquidity shock). However, it is possible that investors will adapt their trading strategies. For example, instead of executing one large order, investors may instead execute several small trades. Recent improvements in technology have fuelled a rise in electronic trading, which has improved the ability to trade small sizes (e.g. O'Hara and Zhou, (2021)). [31]

Finally, it is possible that transparency might impact the primary corporate bond market. Using the introduction of TRACE, Brugler, Comerton-Forde, & Martin, (2021) show that mandated post-trade transparency reduces the cost of issuing corporate bonds, through reduced informational asymmetries. Of course, the same forces that have reduced adverse selection more broadly would also reduce the need for recent secondary market prices as reference points for pricing new issues.

---

[31] In an industry research report, Todorova & Diaz (2023) analyse electronic trading in Europe and conclude that investors use the protocol to trade small tickets in liquid bonds.

References

Adrian, T., Boyarchenko, N., & Shachar, O. (2017). Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics, 89*, 92-109.

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

Back, K., Liu, R., & Teguia, R. (2018). Signaling in OTC Markets: Benefits and Costs of Transparency. *Journal of Financial and Quantitative Analysis, 55*(1), 47-75.

Bao, L., O'Hara, M., & Zhou, X. (2018). The Volcker Rule and corporate bond market making in times of stress. *Journal of Financial Economics, 130*(1), 95-113.

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

Bessembinder, H., & Maxwell, W. (2008). Markets: Transparency and the corporate bond market. *Journal of economic perspectives, 22*(2), 217-234.

Bessembinder, H., Maxwell, W., & Venkataraman, K. (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics, 82*(2), 251-288.

Brugler, J., Comerton-Forde, C., & Martin, J. S. (2021). Secondary Market Transparency and Corporate Bond Issuing Costs. *Review of Finance, 26*(1), 43-77.

Choi, H., Huh, Y., & Shin, S. (2023). Customer Liquidity Provision: Implications for Corporate Bond Transaction Costs. *Management Science, 70*(1).

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

Dick-Nielsen, J., & Rossi, M. (2018). The Cost of Immediacy for Corporate Bonds. *The Review of Financial Studies, 32*(1), 1-41.

Duffie, D., Dworczak, P., & Zhu, H. (2017). Benchmarks in Search Markets. *The Journal of Finance, 72*(5), 1983-2044.

Duffie, D., Gârleanu, N., & Pedersen, L. (2005). Over-the-Counter Markets. *Econometrica, 73*(6), 1815-1847.

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

Feldhütter, P. (2012). The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures. *The Review of Financial Studies, 25*(4), 1155-1206.

Glebkin, S., Yueshen, B., & Shen, J. (2022). Simultaneous Multilateral Search. *The Review of Financial Studies*, 571-614.

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

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.

Goldstein, M., Hotchkiss, E., & Sirri, E. (2006). Transparency and Liquidity: A Controlled Experiment on Corporate Bonds. *The Review of Financial Studies, 20*(2), 235-273.

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

Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.-O., & Zúñiga, D. (2021). Corporate Bond Liquidity during the COVID-19 Crisis. *The Review of Financial Studies, 34*(11), 5352-5401.

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

Meli, J., & Todorova, Z. (2023). Credit ETFs in Mutual Funds and Corporate Bond Liquidity. *Financial Markets, Institutions and 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.

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

Todorova, Z., & Diaz, A. (2023). There's still a role for voice. *Barclays research series*.

Vairo, M., & Dworczak, P. (2023). What Type of Transparency in OTC Markets. *SSRN working paper*. Retrieved from https://ssrn.com/abstract=4431190

Zhu, H. (2012). Finding a good price in opaque over-the-counter markets. *Review of Financial Studies, 25*(4), 1255-1285.

List of Figures

Figure 1: US dealers balance sheet (net positions)

The figure shows net positions of primary dealers in corporate bonds in millions of US dollars.

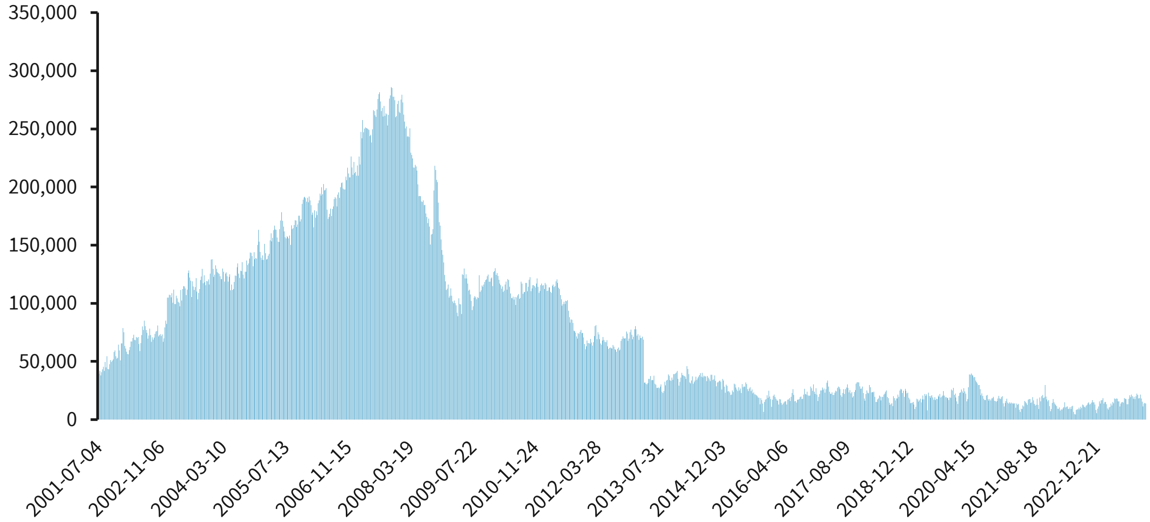

Figure 2: ETF market cap

The figure shows the market cap of the largest IG ETFs in the US (LQD) and in Europe (IEAC).

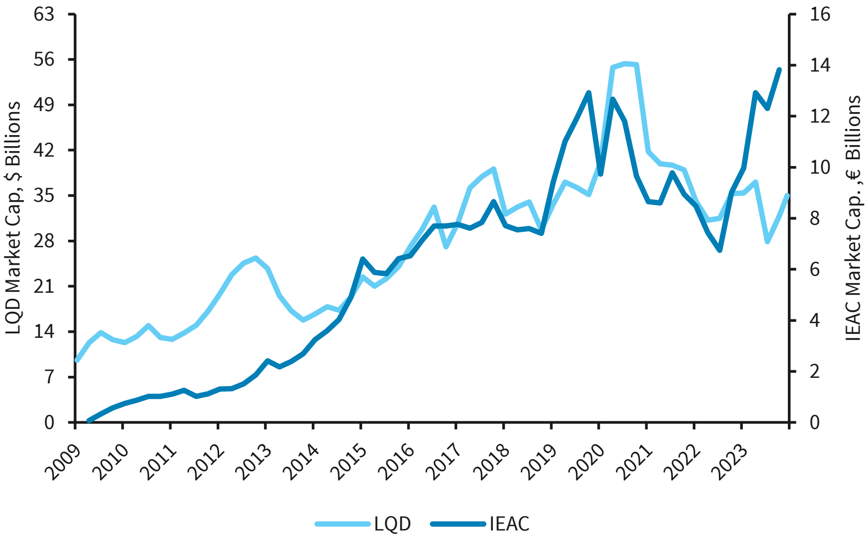

Figure 3: IRC: agency vs. principal trades

The figure shows the average IRC (in bp) for agency and principal trades.

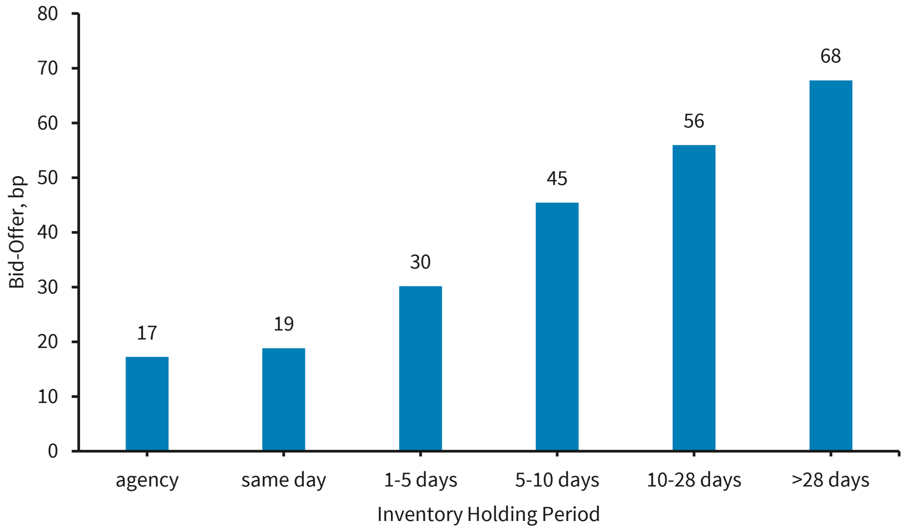

Figure 4: IRC vs. LCS

The figure compared the monthly average weighted IRC and LCS.

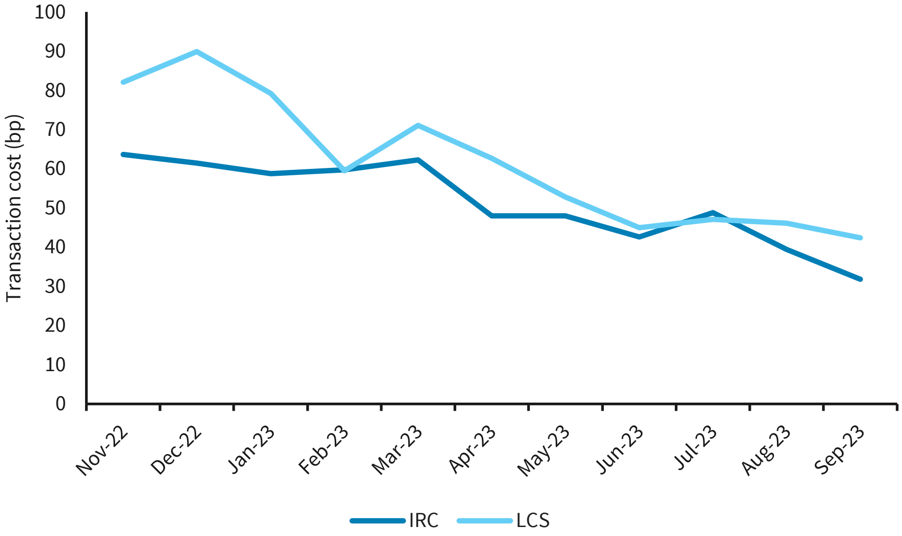

Figure 5: ESMA post-trade reporting rules

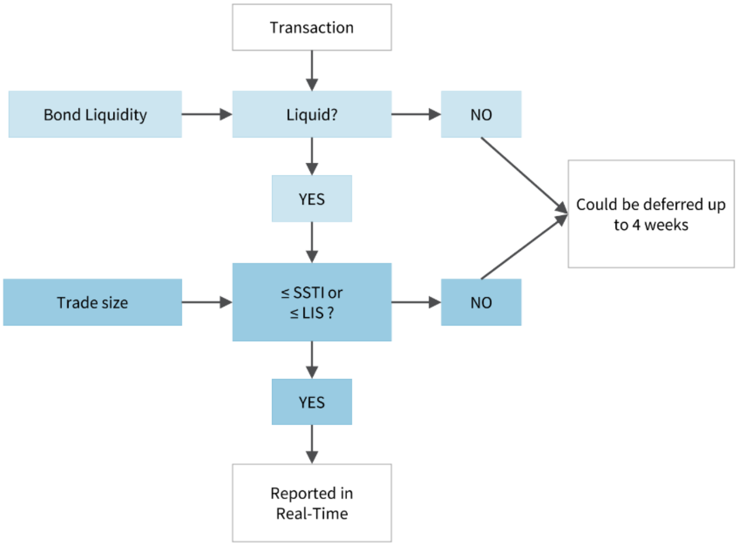

Figure 6: Variation in transactions reporting N.1 – Brexit effect

The figure shows the percentage of trades reported with a delay by size buckets for bonds classified as liquid by the ESMA (dark blue bars). Within each size bucket, we also show what percentage of the reporting variation can be explained by a jurisdiction effect (different liquidity classification and different size threshold in the EU and in the UK), package transaction effect (TPAC) or other sources.

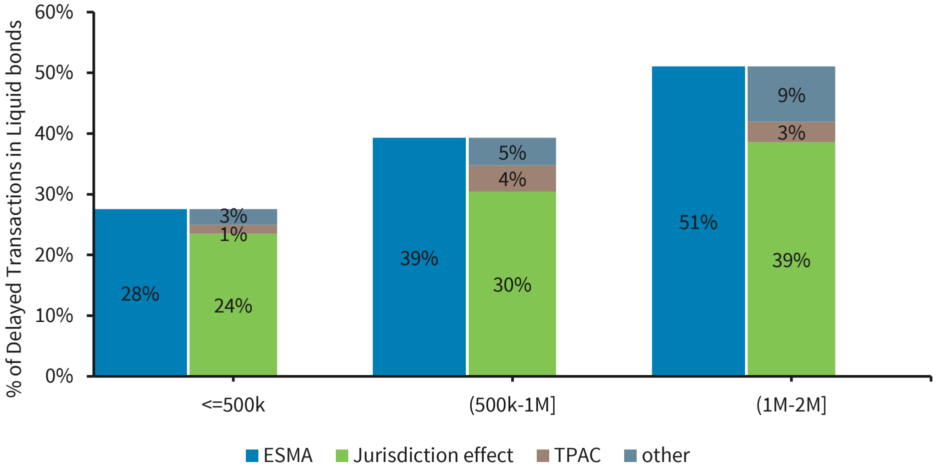

Figure 7: Variation in transparency N.2 – the EU "no publication" event

The figure shows the percentage of transactions reported with transparency (Panel A) and the average age of bonds reported with transparency (Panel B) before and after the "no publication" event between the 16th November 2022 and 16th February 2023.

Panel A: Transaction reported with transparency during the Grey Period

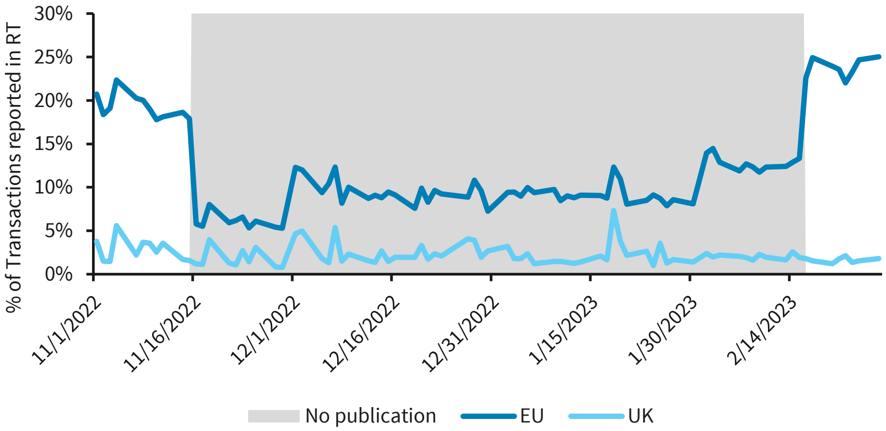

Panel B: Average bond age of transactions reported with transparency

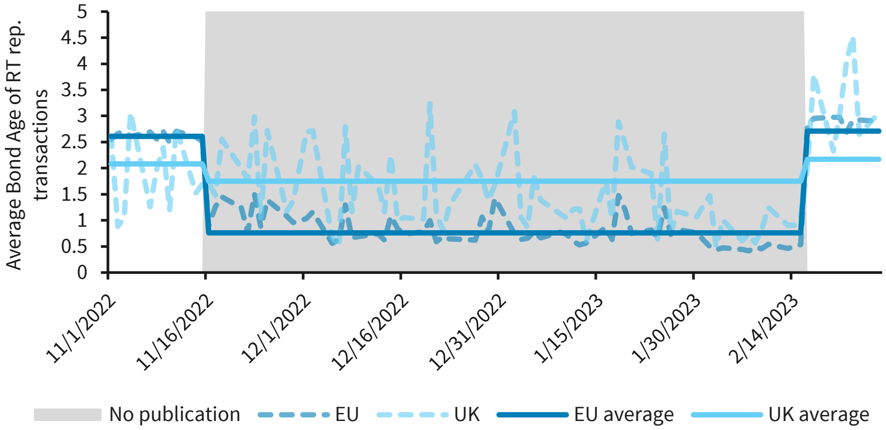

Figure 8: The effect of transparency by trade size

The figure is based on the regression coefficients contained in column (2) of ***Table 4***.

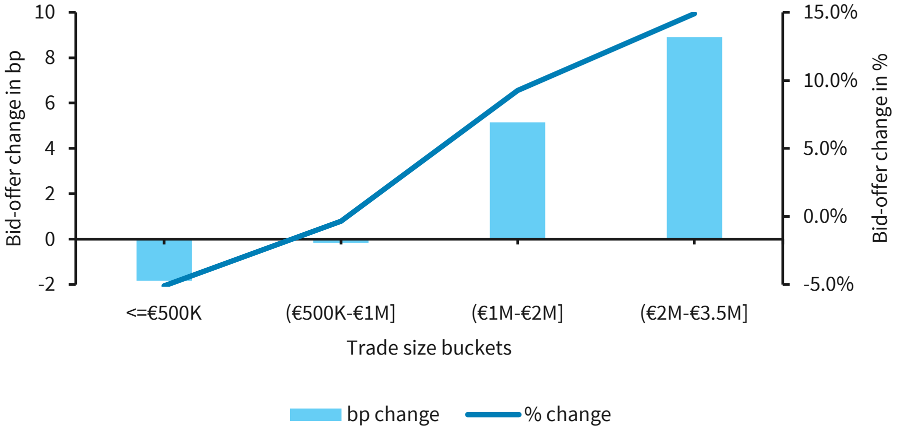

Figure 9: The effect of transparency by bond age

The figure is based on the regression coefficients contained in column (3) of ***Table 4***.

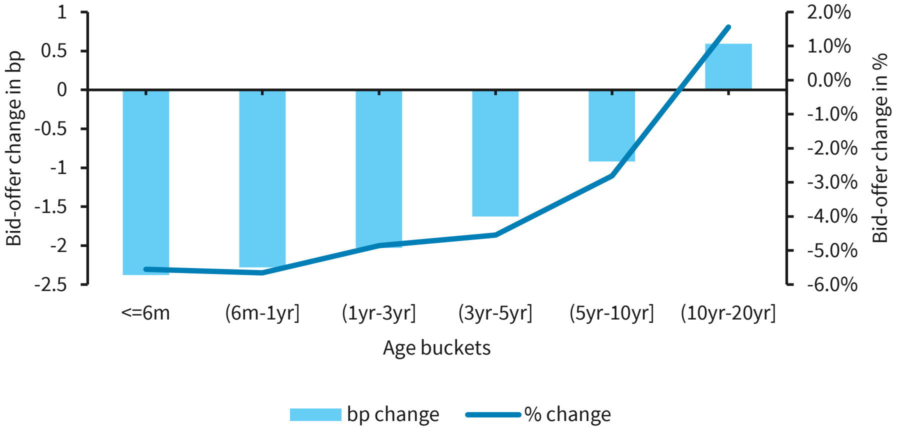

Figure 10: Quasi-natural experiment N.2 – agency trading

The figure shows the percentage of agency trades in treated and control bonds before and after the "no publication" event between the 16th November 2022 and 16th February 2023.

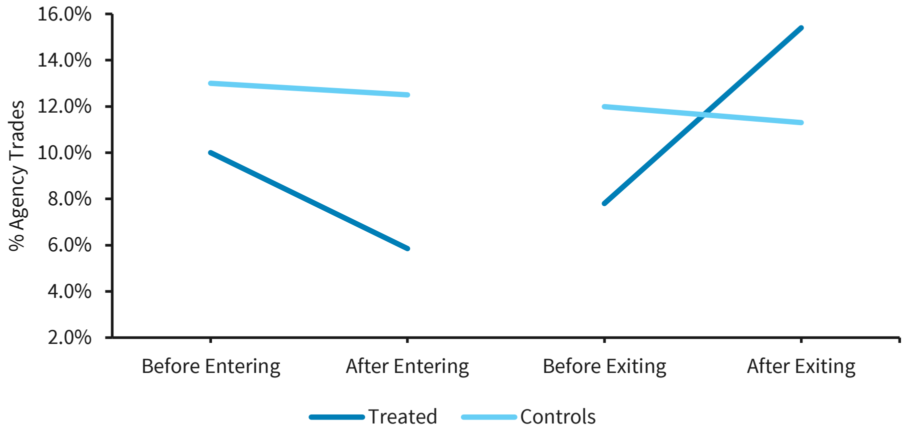

List of Tables

Table 1: A Snapshot of the Data

The table shows an example of the European corporate bond trade data we assemble from the venues we scrape.

Execution date	Reporting date	ISIN	Size	Price	Venue	Jurisdiction	Liquid	Rep. delay	Flag	Flow
29/11/2022 09:38	29/11/2022 09:53	ABC	1.0M	100.63	Bloomberg	EU	YES	5 min	-	Electronic
29/11/2022 10:55	03/01/2023 07:52	ABC	4.0M	101.21	Bloomberg APA	EU	YES	4 weeks	LRGS	Voice
28/04/2023 18:02	30/05/2023 08:44	ABC	300K	99.54	Tradeweb APA	UK	YES	4 weeks	TPAC	Voice
28/04/2023 17:25	30/05/2023 06:56	XYZ	500K	105.54	Tradeweb	UK	NO	4 weeks	ILQD	Electronic

Table 2: Bond Characteristics – EU vs. UK

The table shows summary statistics of the bonds and volumes traded in the EU and the UK.

	(1) EU	(2) UK
Panel A: Trading volume		
Mean round-trip size	€314K	€336K
Total round-trip volume	€138B	€76B
Panel B: Bond characteristics		
Mean Outstanding	€1.028B	€ 1.029B
Mean Age	3.5 years	3.2 years
Mean Maturity	3.6 years	4 years
Unique issuers	704	699
Unique ISINs	2,503	2,468
Bond-round-trip observations	440,963	225,170
Period	Nov-2022 – Sept-2023	

Table 3: IRC – EU vs. UK (Cross-sectional Analysis)

The table compares the bond-level (i.e., cross-sectional mean) of IRC, split by jurisdiction (EU vs. UK) and by type of roundtrip (agency vs. principal)

	Mean IRC, bp			
	Agency trades		Principal trades	
	EU	UK	EU	UK
All bonds	14.5	14.5	43.4	42.5
Bonds with the same liq. classification	14.9	15	43.8	42.7

Table 4: The Effect of Transparency on Transaction Costs

$$\boldsymbol{IRC_{i,j,k,t} = \beta_1 Transparency_{i,j,k,t} + \Gamma X_{i,j,k,t} + \Phi Z_{j,t} + \lambda_j + \delta_k + \gamma_t + \epsilon_{i,j,t}}$$

The table reports regressions at the round-trip level of Imputed Round-trip Cost ($IRC_{i,j,k,t}$) on a transparency dummy ($Transparency_{i,j,k,t}$) and a set of controls. Regressions include the following controls at the round-trip level ($\boldsymbol{X_{i,j,k,t}}$): number of days to close a position, electronic trade dummy and a package transaction dummy. Time-varying controls include ($\boldsymbol{Z_{j,t}}$): bond age (years since issuance), the logarithm of amount outstanding, remaining years to maturity and rating category (AAA is the reference category). Regressions include bond ($\boldsymbol{\lambda_j}$), jurisdiction ($\boldsymbol{\delta_k}$) and time (γ_t) fixed effects. Size fixed effects are based on the following trade size buckets: <€500K, (€500K-€1M], (€1M-€2M] and (€2M-€3.5M]. Column (1) is a baseline; column (2) adds interaction terms with trade size buckets; column (3) adds an interaction term with bond age. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp (Principal round-trips)		
	(1) Baseline	(2) Trade size effects	(3) Age effects
Transparency	-1.40*** (-18.72)	-1.83*** (-12.46)	-2.42*** (-12.63)
Transparency × Size Bucket (€500K-€1M]	-	1.66*** (3.97)	1.94*** (4.58)
Transparency × Size Bucket (€1M-€2M]	-	6.97*** (11.47)	7.28*** (11.91)
Transparency × Size Bucket (€2M-€3.5M]	-	10.74*** (7.53)	11.08*** (7.76)
Transparency × Bond Age	-	-	0.20*** (4.84)
Round-trip level controls	YES	YES	YES
Bond-date level controls	YES	YES	YES
Bond FE	YES	YES	YES
Jurisdiction FE	YES	YES	YES
Time FE	YES	YES	YES
Size FE	YES	YES	YES
Round-Trips Observations	629,223	629,223	629,223

Table 5: The Effect of Transparency on Transaction Costs – Robustness

$$\boldsymbol{IRC_{i,j,k,t} = \beta_1 Transparency_{i,j,k,t} + \beta_2 Transparency_{i,j,k,t} \times Size\ Bucket_s + \beta_3 Size\ Bucket_s + \Gamma X_{i,j,k,t} + \Phi Z_{j,t} + \lambda_j + \delta_k + \gamma_t + \epsilon_{i,j,t}}$$

The table reports regressions at the round-trip level of Imputed Round-trip Cost ($IRC_{i,j,k,t}$) on a transparency dummy ($Transparency_{i,j,k,t}$), an interaction term with size buckets ($\boldsymbol{Transparency_{i,j,k,t} \times Size\ Bucket_s}$) and a set of controls. Regressions include the following controls at the round-trip level ($\boldsymbol{X_{i,j,k,t}}$): number of days to close a position, electronic trade dummy and a package transaction dummy. Time-varying controls include ($\boldsymbol{Z_{j,t}}$): the logarithm of amount outstanding, remaining years to maturity, bond age (years since issuance) and rating category (AAA is the reference category). Regressions include bond ($\boldsymbol{\lambda_j}$), jurisdiction ($\boldsymbol{\delta_k}$) and time (γ_t) fixed effects. Size fixed effects are based on the following trade size buckets: <€500K, (€500K-€1M], (€1M-€2M] and (€2M-€3.5M] (trades in the <€500K are the reference category). Column (1) adds date fixed effects for both legs of the round-trip; column (2) adds trading venue fixed effects; column (3) limits the sample to round-trips where both legs are in the same jurisdiction; column (4) includes jurisdiction-date fixed effects; columns (5) and (6) use agency round-trips. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp					
	Principal				Agency	
	(1) Two-way Date FE	(2) Venue FE	(3) Same Jurisdiction	(4) Jurisdiction-Date FE	(5) Baseline	(6) Trade size effects
Transparency	-4.50*** (-32.18)	-2.23*** (-15.37)	-2.36*** (-13.48)	-2.13*** (-14.76)	-0.19 (0.55)	0.04 (0.10)
Transparency × Size Bucket (€500K-€1M]	4.11*** (9.84)	1.81*** (4.31)	1.82*** (3.61)	1.77*** (4.21)	-	-1.74* (-1.83)
Transparency × Size Bucket (€1M-€2M]	9.89*** (16.31)	7.50*** (12.33)	7.62*** (10.46)	7.39*** (12.14)	-	-2.38* (-1.74)
Transparency × Size Bucket (€2M-€3.5M]	14.77*** (10.38)	11.47*** (8.04)	9.97*** (5.83)	11.35*** (7.96)	-	1.94 (0.77)
Round-trip level controls	YES	YES	YES	YES	YES	YES
Bond-date level controls	YES	YES	YES	YES	YES	YES
Bond FE	YES	YES	YES	YES	YES	YES
Jurisdiction FE	YES	YES	YES	NO	YES	YES
Date FE	NO	YES	YES	NO	YES	YES
Two-way Date FE	YES	NO	NO	NO	NO	NO
Size FE	YES	YES	YES	YES	YES	YES
Venue FE	NO	YES	YES	YES	NO	NO
Jurisdiction-Date	NO	NO	NO	YES	NO	NO
Round-Trips Observations	629,223	629,223	387,134	629,223	36,910	36,910

Table 6: Difference-in-Differences – EU transactions during the “no publication” period

$$IRC_{i,j,t} = \alpha + \boldsymbol{\beta_1 Treated_{i,j}} \times \boldsymbol{Post_t} + \beta_2 Treated_j + \beta_3 Post_t + \Gamma X_{i,j,t} + \Phi Z_{j,t} + \epsilon_{i,j,t}$$

The table reports difference-in-differences regressions of Imputed Round-trip Cost ($IRC_{i,j,t}$) executed in the EU on a Treated dummy ($Treated_{i,j}$), Post dummy ($Post_t$) and their interaction term ($\boldsymbol{Treated_{i,j}} \times \boldsymbol{Post_t}$) and a set of controls. Regressions include the following controls at the round-trip level ($\boldsymbol{X_{i,j,t}}$): number of days to close a position, electronic trade dummy and a package transaction dummy. Time-varying controls include ($\boldsymbol{Z_{j,t}}$): the logarithm of amount outstanding, remaining years to maturity, and BBB rating dummy (equal to one if a bond is rated BBB). Regressions include a trade size fixed effect (for trades larger than 1€M) Results in columns (1), (2) and (3) are based on data from the 1st November 2022 to the 30th November 2022; results in columns (4), (5) and (6) are based on data from the 15th Jan 2023 to the 15th Feb 2023. Columns (2) and (5) include a triple interaction term with the trade size dummy ($\boldsymbol{Treated_{i,j}} \times \boldsymbol{Post_t} \times \boldsymbol{Size\ bucket_s}$) and column (3) and (6) include a triple interaction term with the rating dummy ($\boldsymbol{Treated_{i,j}} \times \boldsymbol{Post_t} \times \boldsymbol{BBB\ rating_{jt}}$). T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp					
	Entering the “no publication” period			Exiting the “no publication” period		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	-6.19*** (-2.50)	-6.30*** (-8.56)	-1.87 (-0.56)	3.23*** (2.85)	3.40*** (3.01)	2.18* (1.80)
Treated × Post × Size Bucket (€1M-€2M]	-	-3.51 (-0.89)	-	-	3.33 (0.70)	-
Treated × Post × BBB rating	-	-	-5.66* (-1.75)	-	-	2.78* (1.85)
Treated	1.86*** (3.00)	2.16*** (3.61)	2.16*** (3.61)	0.11 (0.14)	0.06 (0.07)	-0.64 (-0.80)
Post	1.10 (0.86	2.25*** (5.23)	2.25*** (5.23)	1.07 (1.49)	0.89 (1.25)	1.00 (1.38)
Round-trip level controls	YES	YES	YES	YES	YES	YES
Bond-date level controls	YES	YES	YES	YES	YES	YES
Size FE	YES	YES	YES	YES	YES	YES
Round-Trips Observations	1 Nov 2022- 30 Nov 2022 3,157 trades			15 Jan 2023-15 March 2023 11,875 trades		

Model Appendix

Formally, the strategy of the seller in our model is a decision rule that dictates its response to the bid it receives from the dealer. We must fully specify decision rule, including the off-equilibrium components, to verify that proposed strategy for the dealer is in fact an equilibrium. Implicit in all of the equilibria we compute is the "optimal" decision rule whereby the seller transacts if doing so (weakly) increases its utility (from the starting point of $v - \Delta$), and chooses the trade protocol that maximizes its utility, with any "tie" going to the principal trade:

a) Principal-only or agency-only bid, of the form $B = B(v)$: the seller transacts if $B(v) \geq v - \Delta$

b) "Menu" bid, of the form of $B(v) = v - K$ with certainty (principal trade) or $B(v) = v - X$ with probability p (agency trade): the seller transacts if $\max(\Delta - K, \Delta - X) \geq 0$, and chooses the principal bid if $\Delta - K \geq p(\Delta - X)$.

This decision rule maximizes both the seller's utility and its probability of trading, and thus reflects the weak preference for trading. There are Nash equilibria that involve different decision rules. For example, consider the rule whereby the seller chooses to transact only when it experiences a large liquidity shock. This is a Nash equilibrium when paired with the dealer bidding $v - \Delta_u$; neither player has an incentive to deviate, conditional on the other's strategy, even if the cost of inventory is low. However, this equilibrium is not trembling hand perfect because the seller is forgoing transactions at the (tremble) bid of $v - \Delta_d$; in other words, it is weakly dominated by the optimal decision rule. In fact, the optimal rule outlined above weakly dominates any alternative decision rule, because it maximizes both utility and the probability of transacting.

Therefore, this is the only valid equilibrium strategy for the seller. This intuition greatly simplifies the proofs of the propositions; we need only verify that the proposed equilibria entail the optimal strategies for the dealer and the buyer, conditional on the strategy of the seller.

Proof of Proposition 1

Given that [5] is satisfied and the optimal seller strategy, the optimal strategy for the dealer is a bid of $v - \varDelta_d$. If it increased its bid the dealer would buy no additional securities but reduce its profits, and at a lower bid it would not trade when the seller experiences the small liquidity shock. Due to transparency, the buyer (if it arrives) purchases securities at the dealer's reservation price. QED

Proof of Proposition 2

When adverse selection is high, the payoffs to the seller and dealer are the same as in Proposition 1, and thus the same logic applies, with the exception that the buyer bids $v_l - c$ and only purchases the low value security.

When adverse selection is low, the buyer bids $v_h - c$, which increases the profits from principal trading the low value security. One immediate implication is that [5] is sufficient to ensure that the dealer only utilizes principal trading. Therefore, the payoffs to the seller are unchanged, and the same equilibrium applies, with the only difference being that the buyer overpays for the low value security (which it is willing to do by virtue of adverse selection being low). QED

Proof of Proposition 3

Given the optimal seller strategy, [6] and [7] imply that the optimal dealer strategy is to offer a menu of agency and principal trading, which maximizes dealer profits and induces separation when K is defined as in [4]. Due to transparency, the buyer purchases any bond in inventory at the dealer's reservation price. QED

Proof of Proposition 4

Given the optimal seller strategy, [11] implies that the dealer offers differential liquidity (it bids "high" for the low value security and "low" for the high value security) and that the buyer pays v_h for any security in inventory. All low value securities trade in principal form, and agency trading is limited to the high value security, when the seller experiences the small liquidity shock. QED

Empirical Appendix

A1. Figures

Figure A 1: Distribution of monthly bond turnover, by bond age

The figure shows the distribution of monthly bond turnover, by bond age. The figure is excerpted from Hyman, J. and Konstantinovsky, V. (2023).

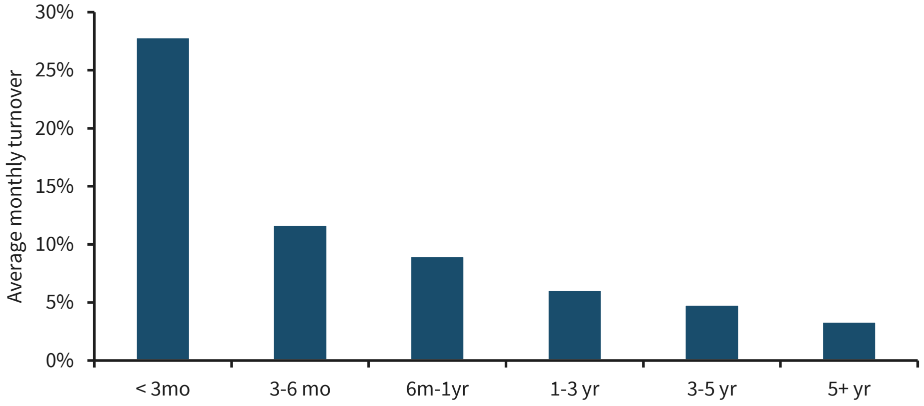

Figure A 2: Distribution of trading activity, by size buckets

The figure plots the percentage distribution of the number of trades and total notional trade by size bucket.

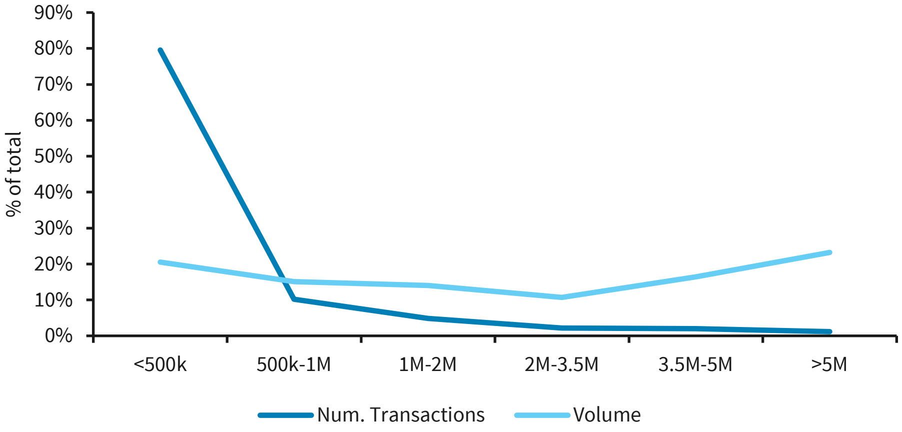

Figure A 3: Parallel trends – agency trading

The figure shows that the parallel trends assumption for treated and control bonds holds.

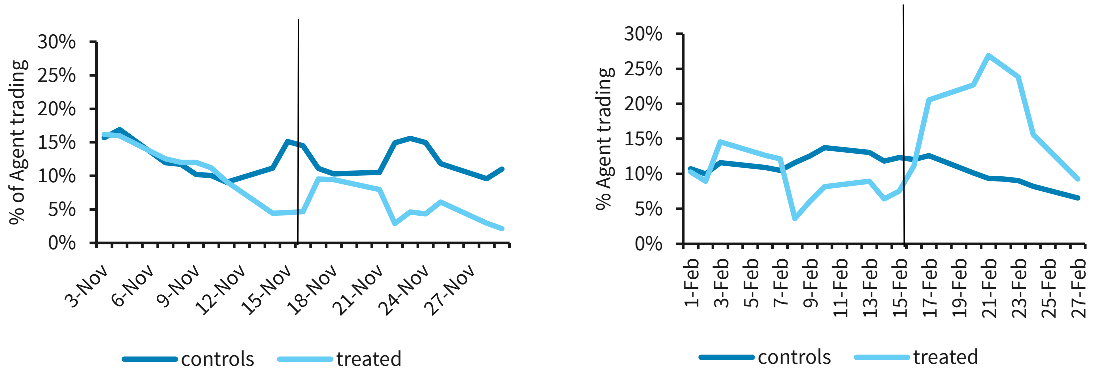

Panel A: Entering the "no publication" period

Panel B: Exiting the "no publication" period

Figure A 4: Parallel trends – bid-offer

The figure shows that the parallel trends assumption for treated and control bonds holds.

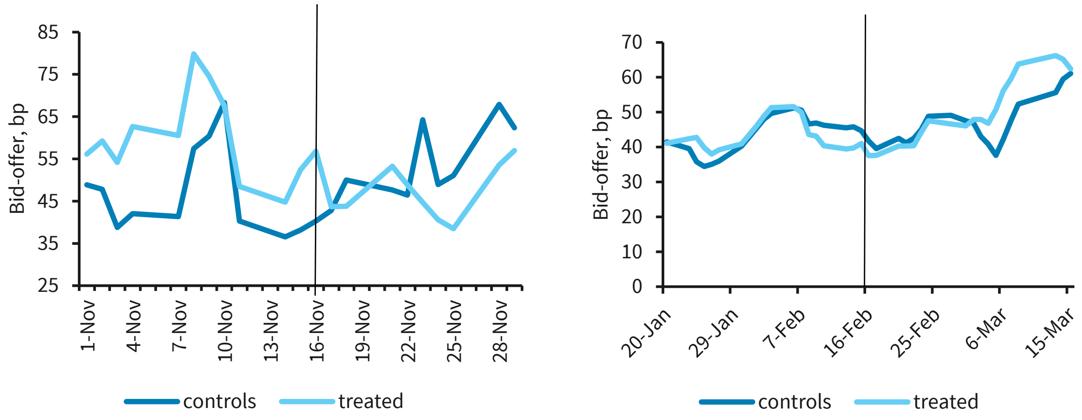

Panel A: Entering the "no publication" period

Panel B: Exiting the "no publication" period

A2. Tables

Table A 1: List of Trading Venues

Jurisdiction	Mifid Entities	Mifid Entities Trading Venue	MIC
EU	Bloomberg	Bloomberg Trading Facility B.V.	BTFE
UK	Bloomberg	Bloomberg Multilateral Trading Facility	BMTF
EU	Bloomberg APA	Bloomberg Data Reporting Services B.V.	BAPE
UK	Bloomberg APA	Bloomberg Data Reporting Services Ltd	BAPA
EU	MarketAxess	MarketAxess NL B.V.	MANL
UK	MarketAxess	MarketAxess Europe MTF	MAEL
EU	TRADEcho	UnaVista TRAEcho B.V.	ECEU
UK	TRADEcho	London Stock Exchange plc	ECHO
EU	Tradeweb	Tradeweb EU B.V.	TWEM
UK	Tradeweb	Tradeweb Europe Limited MTF	TREU
EU	Tradeweb APA	Tradeweb EU B.V.	TWEA
UK	Tradeweb APA	Tradeweb Europe Limited	TREA
EU	TraX	MarketAxess Post-Trade B.V.	TRNL
UK	TraX	Xtrakter Limited	TRAX

Table A 2: Treated vs. controls

The table reports the percentage of trades reported with transparency for treated and control bonds, before and after they enter the “no publication” period, and before and after they exit the “no publication” period.

	% trades reported with transparency in the EU			
	Entering the “no publication”		Exiting the “no-publication”	
	Pre (1-15 Nov 2022)	Post (16-30 Nov 2022)	Pre (15 Jan -15 Feb 2023)	Post (16 Feb -30 Mar 2023)
Controls	88.5%	76.%	81.5%	92.3%
Treated	89.5%	0.0%	0.0%	98.4%