Rethinking Transparency – Evidence from a Quasi-Natural Experiment

Jeffrey Meli †
Barclays, US

Zornitsa Todorova
Barclays, UK

Andrea Diaz
Barclays, UK

13th March 2024

Abstract

We develop a model of market making that incorporates the cost of dealer inventory and the level of adverse selection. With a high cost and low adverse selection, which we argue describes the current corporate bond market, dealers engage in both principal and agent trading. Trade transparency reduces volumes by shifting more trades into the (uncertain) agent protocol, and increases bid-offer of principal trades, particularly for bonds that are hard to "match" in agent trades. To test these predictions, we construct a novel database of euro corporate bond transactions. We exploit exogenous variation in transparency generated by Brexit, and show that transparency decreases transaction costs for small trades but increases transaction costs by 23% for larger and more difficult to match trades. Our results can be used to inform policy makers in light of recent proposals to shorten reporting delays for corporate bond transactions in Europe.

JEL Codes:

Keywords: transparency, liquidity, transaction costs, dealers 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

1. Introduction

In this article we provide both theoretical and empirical evidence that the effect of transparency on over-the-counter (OTC) markets is more nuanced under currently prevailing market conditions. Our results challenge the truism that transparency improves liquidity in OTC markets, based on a series of influential articles that analysed the introduction of the Trade Reporting and Compliance Engine (TRACE) in 2002.[1] Much has changed in the corporate bond market since the inception of TRACE, and two changes in particular could affect the relationship between transparency and liquidity. First, post-crisis financial reforms increased the cost of holding inventory, which can alter the bargaining dynamic between market makers and investors, potentially reducing the incentive to provide immediacy. Second, bond ETFs have grown in size, and trade actively on the secondary market. These instruments provide real-time information about the pricing of corporate risk, which was completely lacking in the pre-TRACE era, when aggregated bond-level transactions were the only source of information about market valuation. ETFs reduce the adverse selection in the bond market, and with it, the potential benefits of transparency.

On the theoretical front, we develop a model of market making in which the cost of inventory, the degree of adverse selection, and the presence of trade transparency jointly determine the equilibrium level of liquidity and the choice of trade protocol. We show that when the cost of inventory is high and adverse selection is low, market makers offer both agent and principal trading options to investors, and that transparency can reduce volumes and increase the bid-offer of principal trades, but only for positions which are difficult to "match" in agent trading. For positions that are easy to match, or when adverse selection is high, transparency improves liquidity.

[1] See Edwards, Harris, & Piwowar (2007); Bessembinder, Maxwell, & Venkataraman (2006); Goldstein, Hotchkiss, & Sirri (2006) along others. These studies are still routinely cited by policy makers considering new regulations designed to increase transparency; for a recent example, see Chapter 6 of the FCA consultation paper.

We test the predictions of our model using a novel dataset of European corporate bond trades, which we construct using public data from a large number of different reporting venues. 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 trade reporting for one quarter. These allow us to isolate the effect of transparency on the transaction costs of principal trades in the same bonds but with different reporting obligations. We find that transparency can reduce liquidity, particularly for positions that are difficult to match, such as large trades. For example, for smaller trades, transparency reduces bid-offer by 3.7%, but for larger trades transparency increases bid-offer by 23%.

Our model incorporates recent structural changes in the corporate bond market. We model an asset seller who faces an uncertain liquidity shock, unobservable to other players, which is the source of potential gains from trade. The potential buyer of the bond faces adverse selection. Differently from previous models, we assume that the buyer makes a take-it-or-leave-it bid to the market maker (as opposed to the reverse, as in Back, Liu, & Teguia (2018)), which reflects the market maker's cost of inventory, and in turn influences the price the market maker quotes the seller. With a low cost of inventory, the market maker engages in principal trading only, and behaves as if the seller faces a small liquidity shock, in order to trade as many bonds as possible. Transparency increases the information available to the buyer, who buys more bonds out of market maker inventory (increasing transaction volume) and/or pays a lower transaction cost, in keeping with the standard results.

However, a high cost of inventory reduces the reservation price of the market maker, who can no longer afford to bid as if the seller faces a low liquidity shock. Instead, it can only make a principal trading bid priced to facilitate a large liquidity shock, reducing the volume and increasing the bid-offer of principal trades. This raises the prospect of agent trading, which can avoid the inventory cost and facilitate transactions when the liquidity shock is low.

We show that the optimal strategy is to offer the seller a menu designed to induce separation: immediacy at a low asset price (i.e., costly principal trading) or uncertain execution at a higher asset price (i.e., cheap agent trading). When the seller faces a large liquidity shock, it chooses immediacy, and vice-versa. However, the availability of agent trading forces the market maker to reduce the bid-offer it charges for principal trades; it must pay above the seller's reservation price or the seller would always prefer the possibility of better execution via an agent trade. The use of agent trading as a means to separation, and the connection between the availability of agent trading and the cost of principal trading, are both new insights, and contributions in their own right.

Low adverse selection creates an additional nuance; absent transparency, the buyer may be willing to overpay for lower value bonds, if the losses from doing so are below its gains from trade. This allows the market maker to offer differential liquidity, whereby it charges low bid-offer for lower quality bonds. Transparency allows the buyer to distinguish between bonds, which alters the balance between trading protocols: some (certain) principal trades in lower value bonds are replaced by (uncertain) agent trading. Those principal trades in lower value bonds that do occur come with elevated transaction costs, particularly when the probability of matching in an agent trade is low, such that the premium paid by the market maker over the reservation price of the seller is low. In other words, under the combination of high costs and low adverse selection, transparency drives more trades into the agent protocol, and the linkage between the option for agent trading on the one hand, and the price of immediacy via a principal trade on the other, can result in differential effects of transparency depending on the probability of executing an agent trade.

This leads to our main hypothesis: the effect of transparency on transaction cost will vary by transaction, depending on the ease with which agent trades can be executed. 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 spread. The easiest way to distinguish between these types of transactions is via the size of the trade. Larger trades are by definition more difficult to match, and thus we expect the effect of transparency to be negative for those trades, but positive for smaller trades.

To test these predictions, we must compare trades with and without transparency executed in the current market environment. To do so, 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 of these trades exists, they are public, and we assemble a comprehensive set of dealer-to-client trades, accumulated from 50 different voice and electronic venues. We use this dataset to estimate the bid-offer spread of round-trip transactions, distinguishing between principal and agent trades.

Our empirical specifications rely on two sources of exogenous variation in trade reporting. The first is linked to Brexit. While the rules for trade reporting are identical in the UK and the EU, they are complex, and vary with the liquidity classification of the bond ("liquid" or "illiquid") and with the size of the trade, where the liquidity classification and size thresholds are calculated using past transactions. Pre-Brexit, all the EU jurisdictions used the same pool of past trades to make these determinations. Post-Brexit, the UK rules are based on data from transactions executed in the UK, and the EU rules are based on data from EU transactions. This results in differential transparency for similarly sized transactions in the same bonds but executed in different jurisdictions.[2] In addition, a data issue in one quarter in 2022 caused a temporary shift in the reporting obligation of some, but not all, bonds traded in the EU, with no effect on reporting in the UK.[3]

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

[3] After the issue was resolved, normal trade reporting resumed.

We exploit this variation in two ways. First, we pool all round-trip transactions in euro denominated corporate bonds executed between November 2022 and September 2023, and examine the relationship between transparency and bid-offer. The overall effect of transparency on principal trades aligns with the prior literature: it generally reduces transaction costs. However, that result is driven by smaller trades, which make up the majority by trade count but a small proportion of total volume. For larger trades, transparency increases realized bid-offer by nearly 23%.

Second, we use the one-time shift in EU reporting in 2022 to perform a series of difference-in-difference regressions. We compare treated bonds, which experienced a change in reporting, to control bonds, which did not. First, we find that the proportion of agent trades in treated bonds declined (vis-à-vis control bonds) when trades were reported with a delay, and then increased once transparency was restored, in keeping with the predictions of our model. Second, we find that the bid-offer spread of treated bonds declined once transparency was reduced, particularly for bonds that are more difficult to match in the agent protocol.

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. Many of these models focus on search costs, as investors navigate opaque markets and adverse selection in an attempt to access liquidity. For example, Duffie, Gârleanu, & Pedersen, (2005) consider the effect of search and bargaining on valuation in OTC markets, and Zhu (2012) studies the effect of a repeat contact with the same buyer as a means to infer asset value. A number of articles examine the effect of transparency on liquidity in search models, including Duffie, Dworczak, & Zhu (2017), and Vairo & Dworczak (2023).[4] The conclusions across these models are mixed. Duffie et al. (2017) find that transparency (in the form of a published

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

benchmark) typically increases liquidity, although the effect on market maker profits 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 & Hotchkiss, (2020)). 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, who faces an infinite cost of inventory, and so never holds inventory. The authors find benefits of transparency. We model a positive, but finite, cost of inventory, and assume that market makers have differential bargaining power: they have all the bargaining power before acquiring an asset (i.e., they can purchase the asset at the seller's reservation price) but none once acquired (i.e., the buyer purchases the asset at the dealer's reservation price).[5] This endogenizes the choice of trade protocol; agent trading only occurs in equilibrium when trading costs are high, when market makers use it to separate more motivated sellers, who pay for immediacy, from less motivated sellers, who choose uncertain but cheap agent trading. We are the first (so far as we are aware) to demonstrate a linkage between these protocols: the availability of agent trading requires lower bid-offer on principal trades, particularly for trades that are easy to match. This complicates the effect of transparency, which can reduce liquidity when adverse selection is low.

[5] Some models (e.g. Duffie et. al. (2005)) include a division of bargaining power between traders, but do not explicitly include a market maker.

We also contribute to the empirical literature that studies how increased transparency in the corporate bond market impacts liquidity. The majority of studies on the subject are 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 papers is that transparency decreases transaction costs for investors 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 the current conditions in the corporate bond market.

2. Model

2.1. Motivation and assumptions

In the two decades since the introduction of TRACE, the corporate bond market has experienced several major changes, which together could have implications for the effect of transparency on liquidity. First, the cost associated with dealer inventory has increased due to post-crisis financial reforms (Dodd-Frank Act, the Volcker Rule and the Basel III framework among others) that raised the capital charges associated with inventory and restricted market-makers' risk-taking behaviour.[6] As a result, inventory shrunk considerably after the crisis.[7] In ***Figure 1*** we show that bank balance sheets grew at an exponential rate in the years running up to the crisis, contracted sharply in late 2008 and continued to decline steadily in the post-

[6] 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.

[7] 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 to see a similar pattern for European market-makers.

crisis years. As of September 2023, market makers' balance sheets are 20 times smaller than in 2006.

A higher cost of inventory could conceivably alter the distribution of bargaining power between market makers and investors, and thus the provision of liquidity. Relatedly, the proportion of transactions done via agent trading has increased (e.g. Goldstein & Hotchkiss, (2020), Choi, Huh, & Shin, (2023)). These transactions are done "on order", meaning that the market maker lines up both sides of the trade before executing. It seems intuitive that the rise of agent trading is linked to higher inventory costs. However, existing models of market making are not well-suited to explore this connection.[8]

Another important change is the rise of corporate bond ETFs, which have extremely high secondary market liquidity (Meli & Todorova (2023)) and provide real-time information about the price of corporate credit risk. When TRACE was introduced these instruments did not exist; as a result, aggregated trade reporting was the only source of information about the overall level of credit spreads. Put differently, the potential gains from transparency were higher then, as potential investors were subject to a greater degree of adverse selection. However, ETF market cap has grown 50 times in the last 20 years (***Figure 2***). This has reduced the degree of adverse selection, potentially reducing the need for transparency, or at least muting its benefits for liquidity.

In order to incorporate these factors, we draw from other models of market making, including Back, Liu, & Teguia, (2018) and Duffie, Dworczak, & Zhu, (2017)), but make several adjustments that allow us to better examine the intersection of inventory cost, trading protocol, and adverse selection. Like the prior models, we assume that a "seller" owns an

[8] 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"). In neither case is there a choice of trading protocol.

asset of uncertain valuation, and experiences a liquidity shock, which is the source of potential gains from trade. The seller approaches a market maker, who makes a take-it-or-leave-it bid to the seller. Upon execution, the market maker approaches a “buyer” to offload the asset. The market maker faces a cost of inventory, which gives it an incentive to sell the asset; bid-offer is the difference between the transaction prices for round trip trades.

The first adjustment we make is to model two sources of uncertainty. We assume that both the size of the liquidity shock and the value of the bond are stochastic.[9] The uncertainty about the liquidity shock means that the market maker has two choices: it can bid “high” (i.e., as if the liquidity shock was low) and do many trades, or bid “low” (i.e., as if the liquidity shock was high) and do fewer trades.

Second, we alter the dynamics between the buyer and the market maker. Other models (e.g. Back, Liu, & Teguia, (2018)) assume that the dealer makes a take-it-or-leave-it offer to the buyer. However, this limits the implications that the cost of inventory has on liquidity provision. In particular, even an infinite inventory cost does not affect the division of bargaining power: the buyer can never earn any of the rents generated by the liquidity shock.[10] It also does not comport with comments from market-makers about how the dynamics with investors shift once a position is in inventory, as opposed to when initiating a position. Therefore, we reverse this assumption: in our model the buyer makes a take-it-or-leave-it bid to the dealer, that is conditioned on any information about the value of the bond, as well as knowledge about the cost of inventory and the resulting effect on the dealer’s

[9] Other models assume that only one of these is stochastic. For example, Back, Liu, & Teguia, (2018) assume that bond value is uncertain, but the size of the liquidity shock is fixed. Duffie, Gârleanu, & Pedersen, (2005) assume that the need to trade varies across investors.

[10] 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).

reservation price. In other words, the cost of inventory determines the division of gains from trade between the dealer and the buyer.

Together, these adjustments imply that liquidity provision depends on the cost of inventory. With a low cost, the dealer captures a larger share of the gains, and thus bids as if the unknown liquidity shock is low, because it prefers small profits from many trades over larger profits from fewer trades. As the inventory cost rises, the market-maker’s reservation price decreases. Hence, the market-maker prefers fewer trades at a lower price, and forgoes some principal trading. Absent some division of the surplus with the buyer, such a trade-off does not exist. Importantly, our assumptions generate realistic implications that align with the decline in volumes and increase in bid-offer that accompanied the greater post-crisis cost of inventory. It also allows for a richer set of implications of transparency, which results in a better informed buyer which can more easily extract rents from the dealer.

Finally, we explicitly allow for both principal and agent trading. We model agent trading as an attempt to pre-arrange both sides of the trade, which does not always succeed. We show that this protocol is not sustainable in equilibrium when the cost of inventory is low. However, when the cost of inventory is high, the market maker uses agent trading to induce separation; the seller chooses agent trading when it faces a small liquidity shock and principal trading when it faces a high liquidity shock. This motivation for agent trading is novel (so far as we are aware), and again aligns with the increased prominence of agent trading post-crisis. It also adds significant nuance to the effect of transparency; we show that the use of agent trading requires tighter bid-offer spreads on principal trades, to prevent sellers from always choosing that protocol. Transparency effects the division of bargaining power between the dealer and the buyer, and thus influences which bonds trade in which protocol. When adverse selection is low, transparency has differential effects on bid-offer depending on the probability of finding a match in the agent trade.

2.2. Primitives

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

"Seller": An investor that owns a bond 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 $\Delta_u > \Delta_d$. We assume that the liquidity shock and the value of the security are uncorrelated.

"Dealer": A market maker in the bond, willing to provide liquidity to the seller. The dealer can either sell the bond once acquired (see below), or hold the bond in its inventory. If it holds the bond, the dealer incurs cost $c > 0$.

"Buyer": An investor who is potentially willing to buy the bond at a price negotiated with the dealer.

We assume that the seller is fully informed about both the value of the bond and the size of its liquidity shock. The dealer observes the value of the bond but not the size of the liquidity shock. The buyer observes neither the value of the bond nor the liquidity shock. However, we assume all players know the distribution of the variables and the cost of inventory. We consider two trading regimes: with and without transparency. Transparency allows the buyer to observe the transaction between the seller and the market maker (akin to TRACE), and thus to infer the value of the bond. Finally, we assume that all players have a weak preference for trading.

Trading

We assume that there are two trading protocols available to the dealer. First, the dealer can make a take-it-or-leave-it offer to the seller. If the seller accepts, the dealer owns the bond in its inventory, and then searches for an interested buyer, who then in turn makes a take-it-or-leave-it offer to the dealer. We label this protocol “principal trading”.

Second, we consider an “agent trading” protocol, in which the dealer can attempt to pre-negotiate a trade by lining up both the buyer and the seller in advance, and execute both sides of the trade simultaneously. In keeping with the bargaining power of the buyer vis-à-vis the dealer (and with the small bid-offer associated with actual agent trades), we assume that the buyer pays a small mark-up, such that the dealer earns a fixed profit of γ on an agent trade. Further, we assume that agent trading is uncertain; the dealer identifies a buyer to match with the seller with probability $p' < 1$. Finally, we assume that the inventory cost is not so high relative to the size of the liquidity shocks that there is no principal trading in equilibrium: $\max[\Delta_d - c,\ q(\Delta_u - c)] > \gamma * p'$.[11] We will determine the conditions under which agent trading occurs in equilibrium, and its effect on volumes and bid-offer.

2.3. Timing and Equilibrium with no buyer

We use a one period model, with the following stages. First, the seller asks the dealer for a bid on the bond; the dealer can make a take-it-or-leave-it bid, offer to accept an order at a given price and execute conditional on a successful attempt to find a match, or offer a “menu” consisting of those options at different prices. Regardless, the dealer knows the value of the bond, and thus $B = B(v)$. If a take-it-or-leave-it offer is accepted, the dealer searches for a prospective buyer, who in turn makes a take-it-or-leave it bid to the dealer, A. When transactions are reported, such that the buyer can observe the price that the dealer paid for the bond, the buyer can differentiate its bid by bond type, $A = A(v)$. Absent transparency, the buyer does not observe the details of the transaction between the seller and the dealer, and

[11] During and after the GFC, the number of agent-only trading firms increased. These typically have little capital and thus face extreme inventory costs. However, overall in the market both principal and agent trading continued, justifying this assumption from a modelling standpoint.

thus cannot distinguish between high and low value bonds, and the buyer's price A is conditioned only on the model primitives. If an agent trading order is accepted, the dealer attempts to find an interested buyer; again, transparency allows the buyer to differentiate between high and low value bonds. For either protocol, we define the bid-ask spread to be the average of $B - A$ for round-trip principal trades.

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

Simple case: No buyer

We first consider a simple case where there is no buyer: the dealer must hold the bond in inventory if it executes a transaction with the seller. Because the dealer knows the value of the bond, we drop the subscripts from v. Due to the discrete nature of the liquidity shock the dealer will make one of two offers for the bond: $v - \Delta_u$ 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 sale price less the value of the bond (including 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 (all bonds trade) and prices that are close to "fair" value. 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 high. The dealer earns profits of $\Delta_d - c$, which is greater than the profit of bidding a lower price and only executing when the liquidity shock is high.

When $c > c'$ we have a low liquidity equilibrium. Transaction volume declines to q, and the discount to fundamental value increases Δ_u (we formalize the bid-offer spread below when we include a buyer). 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 $q(\Delta_u - c)$. This cost threshold delineates between equilibria when we include a buyer, and thus we use it to define low and high cost regimes:

Definition: When [1] is satisfied, inventory cost is low. Otherwise, inventory cost is high.

2.4. Equilibrium with a buyer

Transparency

We now include a buyer, starting with a transparent market. Transparency implies that the buyer can distinguish between bonds based on the posted transaction between the seller and the market maker, and thus we once again drop the subscripts from v. For any bond in inventory, the buyer bids $v - c$, which is the reservation price of the market maker.

The market maker must decide if it will offer only principal trading, or a menu of principal and agent trading:

$$B(v) = v - K \; with \; certainty \; or$$

$$B(v) = \; v - X \; if \; a \; match \; is \; found \qquad [2]$$

The goal of the menu would be to separate high and low liquidity shocks, whereby the seller chooses expensive principal trading when it faces a high liquidity shock, and vice-versa. This is optimal for the market maker if the increased profits from principal trading when the liquidity shock is high outweigh the decline in profits from agent 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 agent trading when it faced a low liquidity shock. However, this implies that the seller would earn positive expected utility from agent 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 agent trading. In other words, the availability of agent trading necessarily reduces the cost of principal trading. The market maker therefore 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:

$$\Delta_u - K = p'(\Delta_u - \Delta_d) \rightarrow$$

$$K = (1 - p')\Delta_u + p'\Delta_d \qquad [3]$$

Low inventory cost

With a low inventory cost, the optimal principal-only trading strategy is to buy all bonds at a price of $v - \Delta_d$, earning profits of $\Delta_d - c$. This is preferable to the menu described in [2] when:

$$\Delta_d - c > q * (K - c) + (1 - q) * \gamma * p' \qquad [4]$$

Substituting for K from [3], and noting that $\Delta_d - c > q * (\Delta_u - c)$, we see that [4] is always true. Therefore, when inventory costs are low, the assumption that agent trading does not dominate principal trading implies that the market maker offers only principal trading, which comports with the observation that agent trading was relatively rare before the GFC.

Lemma 1: When inventory cost is low as in [1], and the profits from agent trading are low enough that agent trading does not dominate principal trading, the equilibrium involves only principal trading.

This leads directly to the equilibrium in the low cost, transparent market.

Proposition 1: The unique low cost, transparent equilibrium is:

a) *The dealer offers $v - \Delta_d$ to the seller, and buys all bonds;*

b) *The buyer buys bonds at $v - c$ from the dealer;*

c) *Total transaction volumes equal 2;*

d) *Realized bid-offer on round trip trades equals $\Delta_d - c$.*

Trading volumes equal 2 because all bonds are bought and sold by the market maker, all on a principal basis. The transaction cost reflects the optimal bid of the market maker (i.e., as if the liquidity shock is low) and the inventory cost.

High inventory cost

When inventory cost is high, the market maker cannot make a principal bid as if the liquidity shock is low. Instead, the optimal principal-only strategy is a bid of $v - \Delta_u$, which only trades with a probability q. Therefore, we first reassess the potential for the menu of trading options outlined in [2] (noting that [3] still applies). The menu is unsustainable in equilibrium when:

$$q * (\Delta_u - c) > q * (K - c) + (1 - q) * \gamma * p' \rightarrow$$

$$(\Delta_u - \Delta_d) > [(1 - q)/q] * \gamma \qquad [5]$$

Equation [5] implies that it is possible to have both agent and principal trading in equilibrium, so long as the gap between the two liquidity shocks is not too large, and/or the probability of the high liquidity shock (which should be rare) is not too high. In other words, if the potential gain from the high liquidity shock is not too large then the market maker prefers to sacrifice some of those gains in an effort to facilitate some trading during a low liquidity shock.

Lemma 2: When inventory cost is high as in [1], the equilibrium can involve both agent and principal trading.

When [5] is satisfied, the equilibrium involves only principal trading and resembles that in Proposition 1, with the exception that volumes fall to $2 * \text{q}$ and bid-offer rises to $\Delta_u - c$. The more interesting case is when [5] is not satisfied:

Proposition 2: The unique high cost, transparent equilibrium with [5] not satisfied is:

a) *The dealer offers the seller a choice of a certain principal trade at $v - K$ or an agent trade at $v - \Delta_d$ (which has success rate p'), for K defined in [3];*
b) *The seller chooses immediacy when it faces a large liquidity shock and the agent protocol when it faces a low liquidity shock;*
c) *All bonds in inventory are sold to the buyer at $v - c$;*
d) *Total transaction volumes equal $2 * q + 2 * (1 - q) * p'$*
e) *Realized bid-offer on round trip principal trades equals $K - c$.*

Here we see that agent 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.

No Transparency

Now we turn to the case without transparency. The immediate implication is that the buyer can no longer differentiate between bonds. Due to the discrete nature of bond type, the buyer will make one of two bids for a bond in dealer inventory: $A = v_h - c$ or $A = v_l - c$.

If the buyer bids $A = v_l - c$ then the dealer will only sell the buyer the low value bond. Alternatively, if the buyer bids $v_h - c$ it will buy both types of bonds from the dealer. While this entails overpaying for the low value bond, it increases the probability of trade, and thus generates a benefit of c across more transactions. The lower bid is optimal if:

$$\theta * c > (1 - \theta) * c + \theta * (c - (v_h - v_l)) \rightarrow \qquad [6]$$

$$c < \left[\frac{\theta}{1-\theta}\right] * (v_h - v_l) \qquad [7]$$

The right hand side of [7] is a measure of the potential for adverse selection. When the probability of the low valuation is small, and/or when the gap between the two bond valuations is small, then the losses associated with indiscriminately buying bonds is low. In other words, with low adverse selection, the buyer is willing to risk overpaying for low value bonds. In contrast, when adverse selection is high, the consequence of overpaying for low value bonds is significant, and the buyer behaves as if all bonds have the low value.

Definition: When [7] is (not) satisfied, adverse selection is (low) high.

Low inventory cost

Like above, with a low cost of inventory, agent trading is unsustainable in equilibrium. This is clearly the case when adverse selection is high. The buyer is only willing to buy bonds at a price of $v_l - c$, and thus agent trading of high quality bonds is impossible. But that forestalls agent trading of low quality bonds, using the same logic as in [4] above. The argument is somewhat more nuanced when adverse selection is low, because the buyer is potentially willing to pay a larger markup for low quality bonds, equal to $\gamma + (v_h - v_l)$.

However, this requires that both types of bonds be available in the agent market. Since we know that the market maker trades all the high quality bonds on a principal basis, only the low quality bonds are available in the agent protocol. Therefore, the buyer does not pay a markup over γ, and we conclude that agent trading does not occur in equilibrium.

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 quality bonds are one-sided. Conversely, when adverse selection is low, removing transparency increases the bid-offer, because the buyer overpays for low quality bonds.

Proposition 3: The unique low cost, no transparency equilibrium is:

a) *The dealer offers $v - \Delta_d$ to the seller, and buys all bonds on a principle basis;*
b) *When adverse selection is high, the buyer buys low value bonds at $v_l - c$, the dealer holds the high value bonds in inventory, total transaction volume is $1 + \theta$, and realized bid-offer on round trip trades equals $\Delta_d - c$;*
c) *When adverse selection is low, the buyer buys all bonds at $v_h - c$, the dealer holds no inventory, total transaction volume is 2, and realized bid-offer on round trip trades equals $\Delta_d - c + \theta * (v_h - v_l)$;*

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

We conclude that when inventory costs are low, transparency always increases liquidity. It either raises volumes, when the degree of adverse selection is high enough to disincentivize trading, or it reduces bid-offer, by obviating the need for buyers to overpay for some bonds. 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 bonds to the dealer regardless. However, the welfare of the buyer increases; it earns rents 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 bonds. Conversely, transparency benefits the buyer because it pays the dealer its reservation price in all transactions.

High inventory cost, low adverse selection

Transparency also increases liquidity when both the cost of inventory and adverse selection are high (see the Appendix for that scenario). However, the implication of transparency can be different when the cost of inventory is high and adverse selection is low.

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

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

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

Note that the original threshold still applies to the high value bond. 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 bonds.

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

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

In order to be an equilibrium, it must be the case that the buyer is still willing to purchase all bonds at the high value even when accounting for the increased proportion of low value bonds. In other words, adverse selection is worse for an inventory cost in this range, such that the constraint in [7] is actually tighter, because we substitute p_l for θ:

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

We now make two assumptions: that the cost of inventory is between the two thresholds and that [10] is satisfied (we consider the other cases in the Appendix). Under these assumptions, the dealer will behave as if its inventory cost is low when presented a low value bond, and as if its inventory cost is high when it is presented a high value bond. This implies that the market maker does not offer agent trading for the low value bond. However, the dealer cannot buy high value bonds when the seller has the small liquidity shock. As above, it instead offers a menu that includes agent trading, to attempt to trade some of these bonds. The buyer purchases all bonds in inventory at a price of $v_h - c$.

Putting these pieces together, we can describe the equilibrium for the high value bond. The dealer behaves as in [2], with K defined by [3]. The seller separates: it chooses immediacy at $v_h - K$ when it faces the high liquidity shock and agent trading at $v_h - \Delta_d$ when it faces the low liquidity shock. This allows us to fully characterize the resulting equilibrium:

Proposition 4: With cost between the two thresholds in [8], low adverse selection (using the threshold based on p_l), and no transparency, the unique equilibrium is:

a) *For the high value bond the dealer offers the seller a choice of a certain principal at $v_h - K$ or an agent trade at $v_h - \Delta_d$ (which has success rate p'), for K defined in [3];*

b) *For the high value bond the seller chooses principal trading when it faces a large liquidity shock and agent trading when it faces a low liquidity shock;*

c) *The dealer buys all low value bonds on a principal basis at* $v_l - \Delta_d$*;*

d) *All bonds in inventory are sold to the buyer at* $v_h - c$*;*

e) *Total transaction volume equals* $2 * [\theta + (1-\theta) * (q + (1-q) * p')]$

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

Total transaction volume reflects the fact that the dealer buys and sells all low value bonds and the high value bonds that are paired with a large liquidity shock, and it (sometimes) matches buyers and sellers of the high value bond when the seller experiences the small liquidity shock. The average bid-offer on principal trades reflects the mix of bonds bought and the differential bid-offer of high and low value bonds.

We now analyse how transparency affects volumes and bid-offer. Total volumes are higher without transparency when:

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

$$1 > q + (1-q) * p' \qquad [11]$$

This is always true (except when $q = 1$). Therefore, with high cost and low adverse selection, transparency changes the mix of trading protocols reduces volumes. Some (certain) principal trades in low value bonds are replaced with (uncertain) agent trades; specifically, when a low value bond is paired with a small liquidity shock, the trade is done on a principal basis absent transparency but on an agent 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 can determine the circumstances in which transparency increases the average bid-offer on principal trades:

$$p_l * [\,\Delta_d + (v_h - v_l)] + (1 - p_l) * K - c < K - c \qquad \rightarrow$$

$$p' < 1 - (v_h - v_l)/(\,\Delta_u - \Delta_d) \qquad [12]$$

According to [12], it is possible that the average bid-offer spread also increases when transparency is introduced. This requires that the probability of an agent match is sufficiently low. The probability of an agent match features in the average principal bid-offer because the discount $K < \Delta_u$ increases 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 the probability of a match is high.

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

This outcome requires both low adverse selection and a cost of inventory that is high but not prohibitive, meaning large enough to deter some principal trading but small enough that the dealer can take advantage of low adverse selection to pool bonds. Under these circumstances, it is still not the case that transparency will reduce liquidity for all trading. It is specifically positions that are more difficult to match via agent trades that suffer from transparency. Positions for which agent trading is relatively easy still benefit from transparency.

2.5. Predictions

This leads to our main testable predictions. We predict that the effect of transparency on transaction costs depends on the size of the transaction, and that transparency changes the mix between principal and agent trades.

H1: Transparency reduces bid-offer spreads for smaller trades but increases it for larger trades.

H2: Transparency increases the proportion of agent trades.

3. Constructing a database of European corporate bond transactions

3.1. Transactions reporting in Europe

In order to test our 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.[12] 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 transactions in euro-denominated investment grade (IG) corporate bonds. 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-

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

makers.[13] 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 reporting business (ie. 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 50 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.).[14] We focus on euro-denominated investment grade (IG) corporate bonds over the period November 2022 – September 2023.[15] For each transaction, we obtain the exact execution and reporting timestamp, 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 both static data (e.g., issuer, sector, issue size) and time-varying bond attributes (e.g., remaining years to maturity, bond age, credit rating (based on a combination of 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 (for more details, refer to ***Table A 3***).

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-

---

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

[14] To account for delays and deferrals in the data, which usually could take four weeks (and in some case longer), we concluded the process of collecting and updating the data as of September 30th 2023 on November 30th 2023.

[15] We collaborated with fintech provider Propellant.digital.

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. (***Figure 3***). 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.[16] Further, while we don't have a precise estimate of the size of the wholesale corporate bond market in Europe, TRACE estimates[17] 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 rate. These tests give us confidence that the dataset we have constructed captures close to 100% of the institutional corporate bond market in Europe.

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

The main drawback of the dataset of trades is that the direction of the trade is not reported under MiFiDII law.[18] Therefore, we measure transaction costs using the imputed round-trip

---

[16] 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.

[17] TRACE explicitly differentiates between dealer-to-client and dealer-to-dealer volumes. Further, Nothing in the existing literature points to systematic differences between the US and Europe in terms of the share of dealer-to-dealer volumes.

[18] 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). Unfortunately, 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. However, these would produce noisy estimates when applied at the transaction level.

cost (IRC) (Feldhütter, (2012); Kargar, et al., (2021)). This measure can be applied at the trade level and can be computed independently of the direction of trade.

To construct the IRC, we first identify pairs of round-trip trades. A round-trip trade consists of two trades in the same bond with the same trade size that are executed as close as possible to each other but have different prices (for an example how our methodology works, refer to ***Table A 2***). 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 buy of another investor from the same market-maker, or vice versa. Then, for each round-trip trade in our list, 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. To ensure that our results are not polluted by extreme values, we remove *IRC* values above the 95th percentile of the distribution.

We explicitly differentiate between agent and principal round-trips. Following the literature (Kargar, et al., (2021)), we identify agent round-trips as trades executed within 15 minutes.[19] Our final sample contains c.666K observations (roundtrips), of which c.630K are principal round-trips and c.36K are agent round-trips.

Robustness

To check the robustness of the *IRC* methodology we do two different tests – (1) we compare the transaction costs of agent and principal round-trips, and (2) compare the

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

transaction cost estimates produced by the *IRC* methodology to the Barclays Liquidity Cost Score (LCS).

Kargar, et al., (2021) find that because market-makers do not use their balance sheet when they intermediate agent trades, agent trades typically cost less than principal trades. We confirm this for our sample – agent round-trips cost on average 17.2bp compared to 38.3bp for principal round-trips (***Figure 4***). 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). We find that *IRC* closely tracks LCS (***Figure 5***).

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 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 6***).[20] 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.[21] Each year, ESMA also sets two global trade size thresholds.[22] 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. 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

[20] 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.

[21] 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.

[22] 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.

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*".[23] 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.[24]

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.

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

[24] Note that 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. Package trades and portfolio trades are reported and treated differently by the regulator.

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.[25]

To demonstrate, in ***Figure 7*** 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

[25] 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.

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 transactions smaller than €500K are reported with a delay (dark blue bars in ***Figure 7***). 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. As discussed previously, another (albeit small) portion of the variation can be explained by package transactions (TPAC). Finally, 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…*".[26] Similarly, in order to be eligible to trade with the EU entity of Tradeweb, an investor must be

[26] Tradeweb UK's Rulebook.

"*...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...*".[27] 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 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.[28] In accordance with the MiFiDII playbook, all bonds for which no liquidity assessment had been published were deemed illiquid[29] 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.

[27] Tradeweb EU's Rulebook.
[28] The press release can be found here.
[29] This is in line with Q&A 10 of section 4 of the MiFID II transparency Q&As.

In the EU, the number of transactions reported with transparency decreased sharply exactly on the 16th November 2022 and subsequently recovered exactly on the 16th February 2023, when the next regular publication period began and the new reporting rules applied (Panel A, ***Figure 8***). In the UK on the other hand, the number of transactions reported with transparency remained 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 consistent with ESMA's guidance regarding newly issued bonds (Panel B, ***Figure 8***). 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 (***Figure A 1*** in the Appendix).

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 (roughly the difference between the light blue and dark blue lines in Panel B of ***Figure 8***, a subset of the bonds aged between 6 months and 2 years).
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 – first, as treated bonds enter the "no publication" period, and second, as treated bonds 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 reduces bid-offer for smaller trades but increases it for larger trades

5.1. Summary statistics and selection bias

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 reported 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. The most important control in this group is bond age. IG bonds typically trade very frequently shortly after they are issued, after which their liquidity sharply declines (***Figure A 2*** in the Appendix).

Finally, we also include bond (λ_j), jurisdiction (δ_k) and time (γ_t) fixed effects to account for any (potentially unobservable) factors that could affect our results. We estimate the model separately for agent and for principal round-trips.

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***) but it has no statistically significant effect on agent trades (column (2) 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 by size. 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.

To examine how the effect of transparency varies across trade sizes, we augment ***Model 1*** by interacting the $Transparency_{i,j,k,t}$ dummy with $Size\ Bucket_s$:

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

where $Size\ Bucket_s$ is one of the following four buckets: ≤€500K, (€500K-€1M], (€1M-€2M] and (€2M-M3.5M]. We omit the ≤€500K category, which is our reference size 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. Similarly, as before, we estimate **Model 2** separately for principal and agent trades.

We find that 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 (1), ***Table 5***). However, transparency has no effect on the transaction costs of agent trades, regardless of trade size (column (2), ***Table 5***). Both of these results are aligned with the predictions of our theoretical model.

Using the regression coefficients in ***Table 5***, in ***Figure 9*** 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.23% increase.

One feature of the corporate bond market is that the number of trades and the value of trades are extremely unequally 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 3*** in the Appendix). Weighing the effect of transparency for each trade size bucket shown on ***Figure 9*** 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.

5.4. Robustness

Our dataset collects data from 50 trading venues (***Table A 1***). Although we have a rigorous regression specification which includes bond, date and jurisdiction fixed effects, it is still possible that we have omitted some (potentially unobservable) trading venue-related factor. Our results could be biased if for whatever reason some venues have both higher transaction costs and are both more likely to report trades with a delay. In columns (1) and (2) of ***Table 6*** we include venue fixed effects and find very similar results to the baseline in ***Table 5***. Further, we obtain similar results if we include jurisdiction-date fixed effects, which control for market events which affected a specific jurisdiction on a given day ((3) and (4)). Finally, limiting our sample to round-trips where both legs are in the same jurisdiction does not affect our estimates ((columns (5) and (6)).

6. "No publication" quasi-natural experiment

6.1. Treated and controls

The "no publication" event exogenously caused the reporting schedule of a subset of the bonds in our sample to change, 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.[30] 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, as they were not classified as recently issued by ESMA.
- **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 7*** 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

[30] 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.

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. Agent trading

We compute the proportion of agent trades for treated and control bonds shortly before and after they enter and exit the “no publication” period (***Figure 10***). The proportion of agent trading for control bonds remains unchanged both as they enter and exit the quasi-natural experiment window (***Figure A 4*** in the Appendix verifies the parallel trends assumption with daily data). However, as treated bonds change their reporting from transparency to no transparency (i.e., as they enter the “no publication” event) the proportion of agent 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. As treated bonds change their reporting from no transparency to transparency, their proportion of agent trading increases from 8% to 15.4%. These results align with the predictions of our theoretical model.

6.3. Difference-in-differences 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 takes a treated bond before and after the treatment 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:

$$IRC_{i,j,t} = \alpha + \beta_1 Treated_{i,j} \times Post_t + \beta_2 Treated_j + \beta_3 Post_t + \Gamma X_{i,j,t} + \Phi Z_{j,t} + \epsilon_{i,j,t} \quad \textbf{(Model 3)}$$

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 – (1) as bonds enter and (2) as bonds exit the "no publication" window (***Table 8***).

The DID estimate relies on two assumptions – (1) that treated and controls are similar, and (2) 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 5*** 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 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 increases by 3 bp ($\beta_1 > 0$). Both results are statistically significant, and overall 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 agent trade:

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

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 agent 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 typically don't trade frequently limits the sample available for inference.

7. Discussion and policy implications

7.1. Forthcoming Changes to the Post-Trade Reporting Rules in the EU and in the UK

In a mission 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.[31] Although some specific provisions and technical details differ slightly, the overarching goal of the review in both jurisdictions is to implement faster disclosures for corporate bond trades and establish a consolidated tape (CT), 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.[32] 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 and leaned on the existing literature based on TRACE data, which concludes that transparency unequivocally improves liquidity for all corporate bonds.[33] 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. In particular, we highlight two market developments since the introduction of TRACE in 2002 – (1) the increase in the

---

[31] 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 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.

[32] 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.

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

cost of market-makers inventory in the aftermath of the GFC and (2) the rise of corporate bond ETFs, which provide a convenient tool for price discovery and which can help mitigate the negative impacts of adverse selection.

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 – maximise transparency subject to preserving liquidity. Our results could provide a helpful starting point; small trades benefit from transparency, but the new framework should make provisions to protect larger trades 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.

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. Transparency can level the playing field and support liquidity for retail investors.

On the other hand, transparency decreases liquidity for the largest and most difficult to match trades. Institutional investors, such as mutual funds, which are more likely to trade larger tickets, might face higher transaction costs. This might be a particularly binding constraint when they are forced to sell quickly and require immediacy – e.g., when investors face a sizeable outflow, which is modelled as a high liquidity shock in our theoretical framework. However, it is also possible that investors adapt, and instead of executing one large order, they execute several small trades. For example, improvements in technology have fuelled the rise in electronic trading, which has made it possible for investors to quickly and efficiently trade in small sizes (e.g. O'Hara and Zhou, (2021)). [34]

Increased transparency could also generate more trading volumes, despite the higher bid-offer associated with large trades, both by improving investor confidence by making it easier to demonstrate "best execution" and by attracting international investors, which would otherwise have been deterred by the lack of pricing data. Transparency could also help contain the volatility of prices, particularly during periods of crisis when the costs and risks associated with adverse selection increase. These are interesting avenues for future research.

Finally, it is possible that secondary market transparency might impact primary corporate bond market through the cost of capital channel. Using the introduction of TRACE, Brugler, Comerton-Forde, & Martin, (2021) show that mandated post-trade transparency reduces the cost of issuing corporate bonds. The mechanism works through reducing informational asymmetries. The authors argue that in bond markets, recent secondary market prices of other comparable bonds are important reference points for pricing new issues, and hence, when a larger fraction of trades in comparable bonds are made transparent, new issue pricing improves. It would be interesting to explore if any causal links exist between the primary and secondary corporate bond market under the current market conditions.

[34] 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.

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.

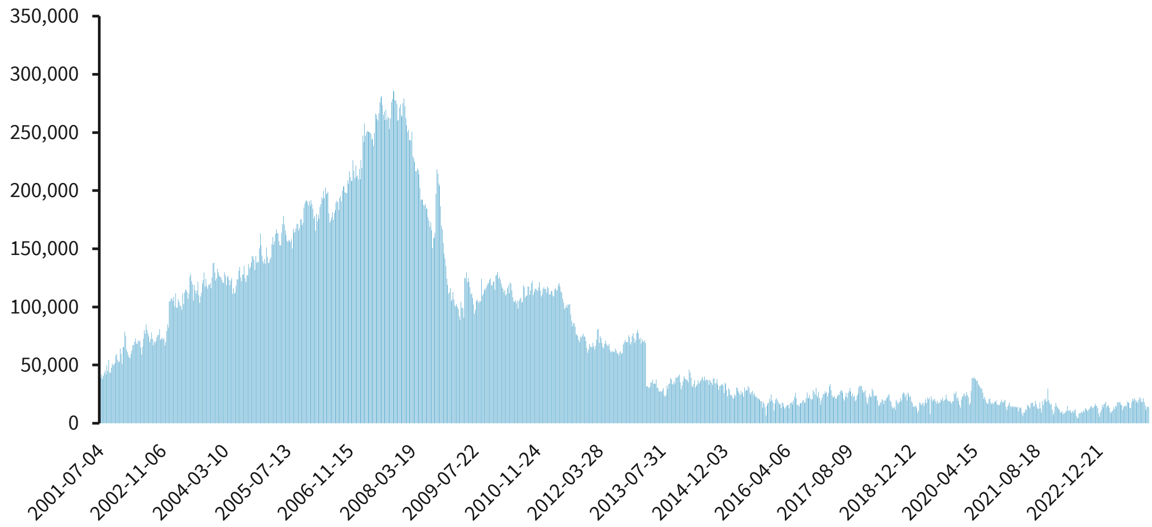

Figure 2: ETFs market cap

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

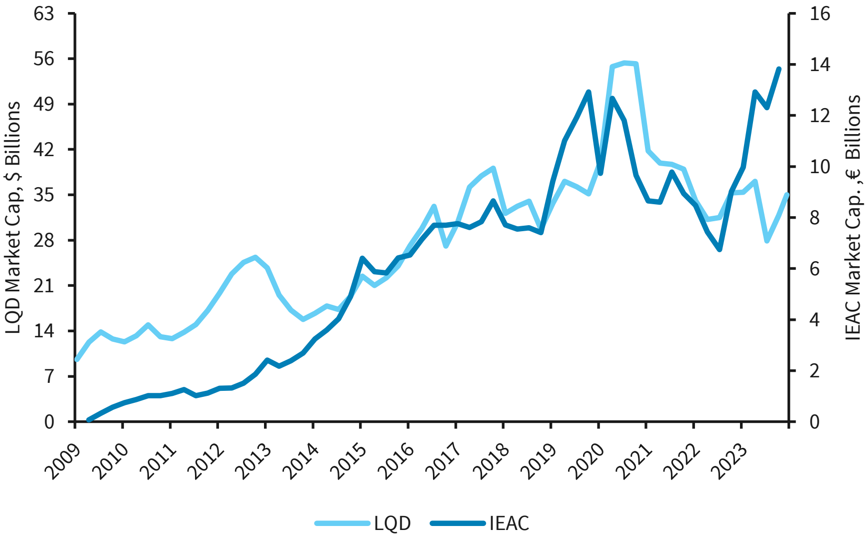

Figure 3: Data quality and matching

The figure shows the monthly percentage of Barclays RFQs (by count and by volume) which we were able to match to our database of corporate bond trades.

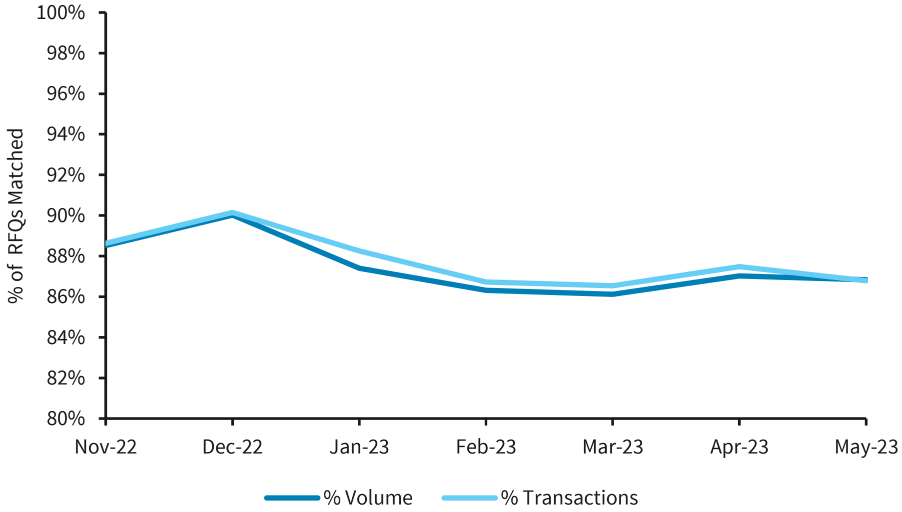

Figure 4: IRC: agent vs. principal trades

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

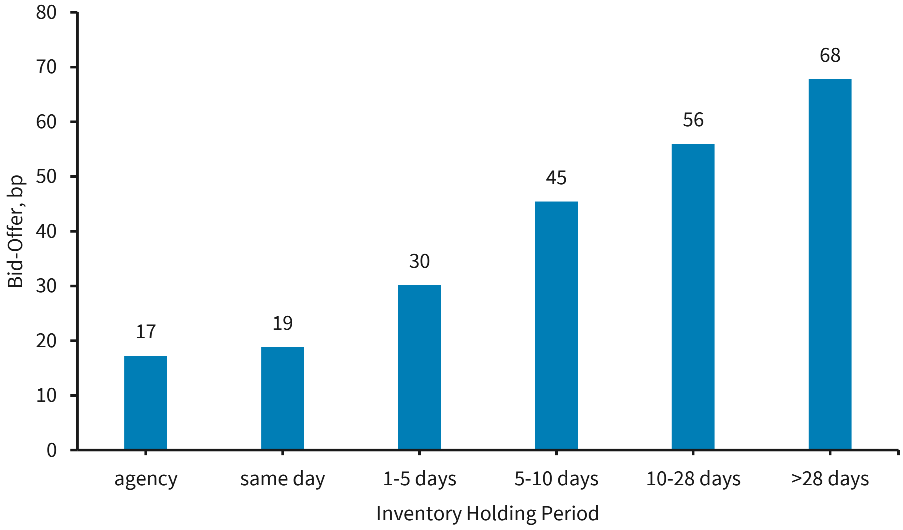

Figure 5: IRC vs. LCS

The figure compared the monthly average weighted IRC and LCS.

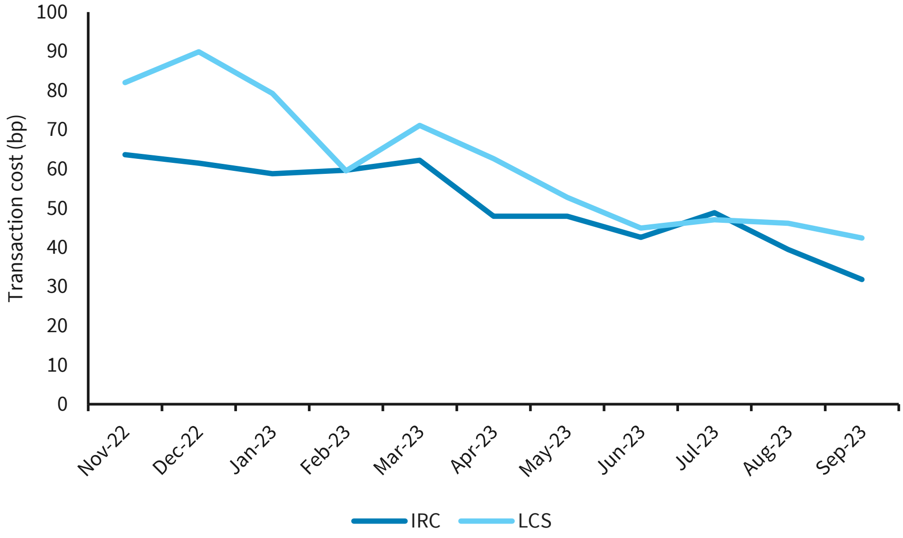

Figure 6: Post-trade reporting rules

The figure shows ESMA's post-trade reporting rules.

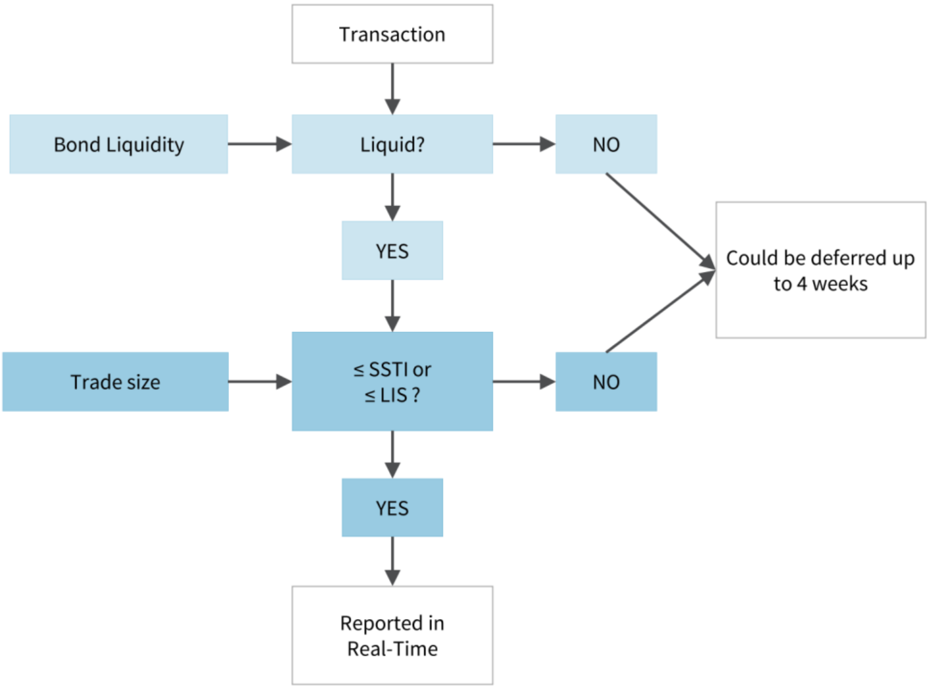

Figure 7: Variation in transactions reporting N.1 – Brexit effect

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.

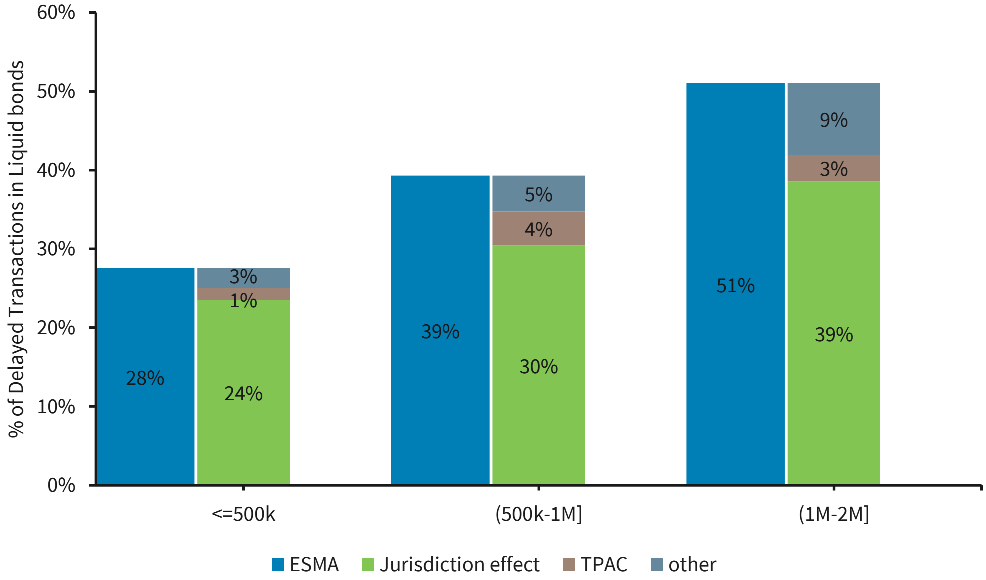

Figure 8: Variation in transactions reporting N.2 – shutting down transparency in the EU

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

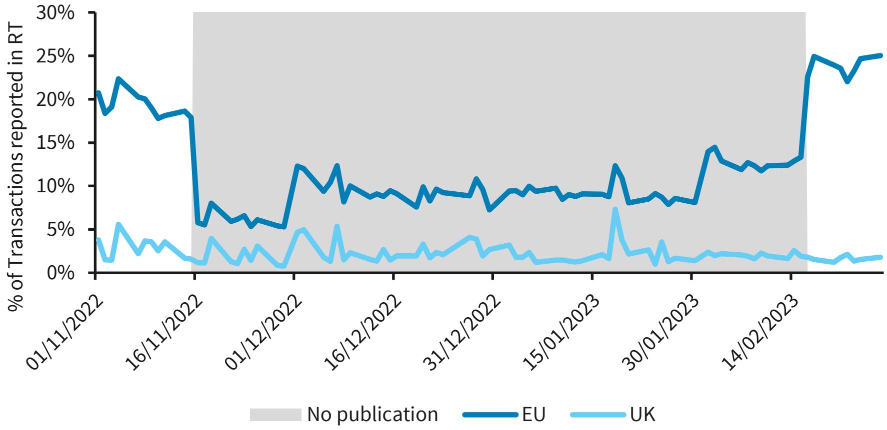

Panel B: Average bond age of transactions reported with transparency during the Grey Period

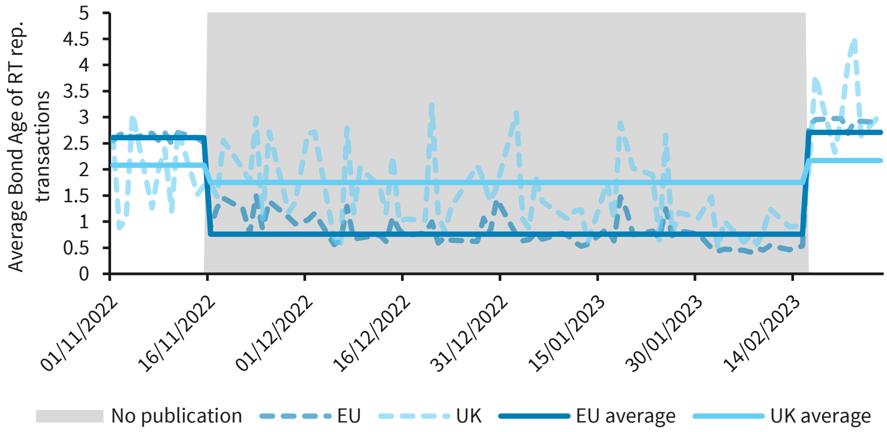

Figure 9: The effect of transparency by trade size buckets

The figure plots the regression coefficients contained in ***Table 5***.

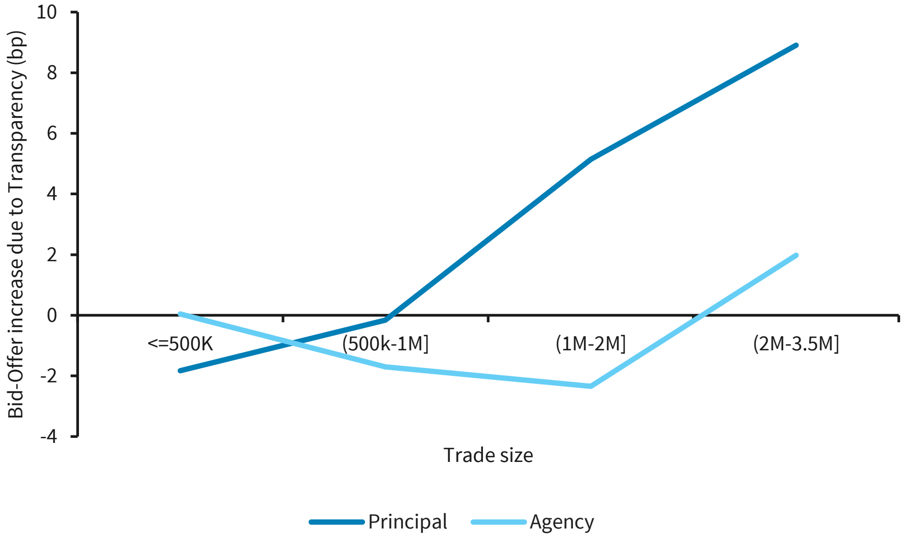

Figure 10: Quasi-natural experiment N.2 – agent trading

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

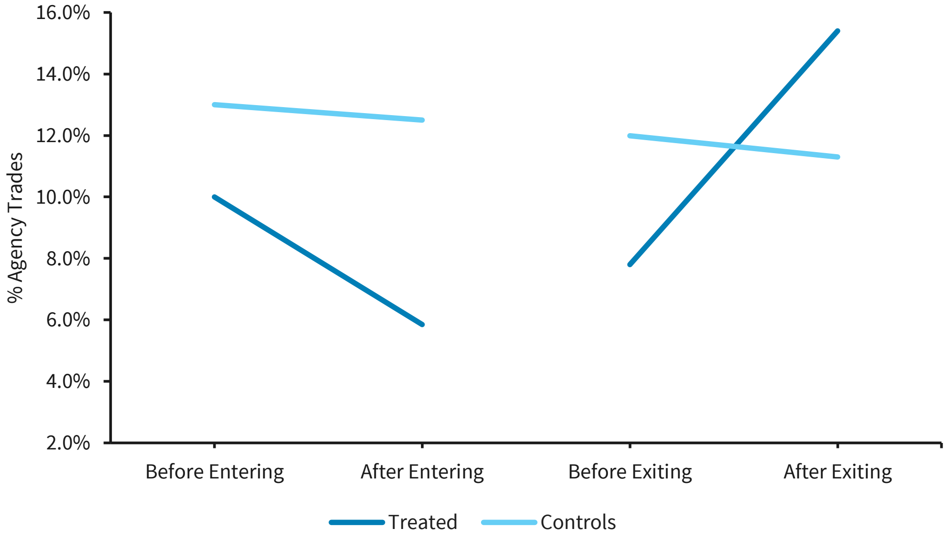

List of Tables

Table 1: A Snapshot of the Data

The table shows an example of the data.

Execution date	Reporting date	ISIN	Size	Price	Venue	Jurisdi ction	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 (ie. cross-sectional mean) of IRC, split by jurisdiction (EU vs. UK) and by type of roundtrip (EU vs. UK)

	Mean IRC, bp			
	Agent 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

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

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 ($X_{i,j,k,t}$): number of days to close a position, electronic trade dummy and a package transaction dummy. Time-varying controls include ($Z_{j,t}$): 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 (λ_j), jurisdiction (δ_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]. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp	
	(1) Principal	(2) Agent
Transparency	-1.40*** (-18.72)	-0.19 (0.55)
Round-trip level controls	YES	YES
Bond-date level controls	YES	YES
Bond FE	YES	YES
Jurisdiction FE	YES	YES
Time FE	YES	YES
Size FE	YES	YES
Round-Trips Observations	629,223	36,910

Table 5: The Effect of Transparency on Transaction Costs – By Trade Size

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

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) uses principal trades; column (2) uses agent trades. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp	
	(1) Principal	(2) Agent
Transparency	-1.83*** (-12.46)	0.04 (0.10)
Transparency × Size Bucket (€500K-€1M]	1.66*** (3.97)	-1.74* (-1.83)
Transparency × Size Bucket (€1M-€2M]	6.97*** (11.47)	-2.38* (-1.74)
Transparency × Size Bucket (€2M-€3.5M]	10.74*** (7.53)	1.94 (0.77)
Round-trip level controls	YES	YES
Bond-date level controls	YES	YES
Bond FE	YES	YES
Jurisdiction FE	YES	YES
Date FE	YES	YES
Size FE	YES	YES
Round-Trips Observations	629,223	36,910

Table 6: The Effect of Transparency on Transaction Costs – Robustness

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

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). Columns (1) and (2) add trading venue fixed effects; columns (3) and (4) limit the sample to round-trips where both legs are in the same jurisdiction; columns (5) and (6) include jurisdiction-date fixed effects. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp					
	Venue		Same jurisdiction		Jurisdiction-Date FE	
	(1) Principal	(2) Agent	(3) Principal	(4) Agent	(5) Principal	(6) Agent
Transparency	-2.23*** (-15.37)	0.16 (-2.94)	-2.36*** (-13.48)	0.63 (1.50)	-2.13*** (-14.76)	1.40*** (3.89)
Transparency × Size Bucket (€500K-€1M]	1.81*** (4.31)	-1.97** (-2.10)	1.82*** (3.61)	-3.16*** (-3.08)	1.77*** (4.21)	-1.93** (-2.07)
Transparency × Size Bucket (€1M-€2M]	7.50*** (12.33)	-2.73** (2.02)	7.62*** (10.46)	-4.24*** (-2.85)	7.39*** (12.14)	-2.84** (-2.10)
Transparency × Size Bucket (€2M-€3.5M]	11.47*** (8.04)	0.22 (0.09)	9.97*** (5.83)	-2.00 (-0.76)	11.35*** (7.96)	-0.27 (-0.11)
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	YES	NO	NO
Date FE	YES	YES	YES	YES	NO	NO
Size FE	YES	YES	YES	YES	YES	YES
Venue FE	YES	YES	YES	YES	YES	YES
Jurisdiction-Date	NO	NO	NO	NO	YES	YES
Round-Trips Observations	629,223	36,910	387,134	27,356	629,223	36,910

Table 7: 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		Exiting	
	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%

Table 8: 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
High cost and high adverse selection

When adverse selection is high, the buyer behaves as if all bonds have the low value (this applies to both principle and agent trading because the dealer cannot commit to only offering high value bonds in the agent protocol).

Therefore, we first assess if the strategy in [4] is tenable for the low value bond. We compute the lowest price the dealer is willing to accept from the buyer such that the dealer prefers to make the bid described in [4], versus the alternative of a principle trading bid at a price of $v_l - \Delta_u$ (which is accepted only when the liquidity shock is large):

$$q(\Delta_u - c) \leq \; q(K - c) + (1 - q) * p' * (\Delta_d - X) \qquad [6]$$

where X is the discount to fair value that the dealer agrees to with the buyer. We substitute for K using equality in [5] and solve for the constraint on X:

$$X \leq c' \qquad [7]$$

Equation [7] implies that, so long as the dealer is able to negotiate a price with the buyer that mimics the price the buyer would have paid were inventory costs actually “low”, then the dealer prefers the menu described in [4] over a strategy of principle trading only at a price of $v_l - \Delta_u$. Note that the buyer utility is increasing so long as $X > 0$, and thus such a bargain is sustainable in equilibrium. In keeping with the assumption that the buyer extracts maximal rents from the dealer (and the market observation that agent trading costs are low) we assume that the dealer sells at the lowest acceptable price, which implies that equality obtains in [7].

Putting these pieces together, we can describe the equilibrium for the low value bond. The dealer behaves as in [4], with K defined with equality in [5]. The seller separates: it accepts immediacy at $v_l - K$ when it faces the high liquidity shock and agent trading at $v_l - \Delta_d$ when it faces the low liquidity shock. The buyer purchases bonds out of the dealer inventory at $v_l - c$, and (when a match is found) engages in agent trading at $v_l - c'$.

We now consider the high value bond. First, note that the buyer does not changes its strategy, which is optimal by virtue of the high level of adverse selection. Knowing this, the dealer has two choices: it can offer the seller only a principle trading option, priced such that it is only accepted when the liquidity shock is high, or it can offer a similar menu that includes agent trading, in an attempt to transact in some situations with a low liquidity shock.

In the first option, the dealer would bid $v_h - \Delta_u$, and realize profit of $\Delta_u - c$. Note that the dealer would not sell the bond to the buyer, whose bid is below the dealer's reservation price. Therefore, expected profits are $q * (\Delta_u - c)$.

In the second option, the dealer would offer the same menu in [4], but substituting v_h for v_l. This is the lowest bid the dealer can make that induces separation. If the seller chooses the principle trade, the dealer will hold the bond in inventory, resulting in expected profits of $q * (K - c)$. If the seller chooses the agent trade, the dealer expects profits of $(1 - q) * p' * (\Delta_d - c' - (v_h - v_l))$, where the last term reflects the fact that the buyer's price presumes that the bond has the low value. The dealer will choose the first option when:

$$q * (\Delta_u - c) > q * (K - c) + (1 - q) * p' * (\Delta_d - c' - (v_h - v_l)) \qquad [8]$$

When we substitute for K from [5] and simplify we find that this equality always holds (as long as $v_h > v_l$); therefore we conclude that the dealer always prefers principle trading only.

This allows us to fully characterize the resulting equilibrium:

Lemma 4: With high cost, no transparency, and high adverse selection, the unique equilibrium is:

g) For the low value bond the dealer offers the seller a choice of $v_l - K$ with certainty or the ability to "work an order" for $v_l - \Delta_d$ (which has success rate p'), for K defined in [5] above;

h) For the high value bond the dealer offers the seller $v_h - \Delta_u$

i) For the low value bond the seller chooses immediacy when it faces a large liquidity shock and agent trading when it faces a low liquidity shock;

j) For the high value bond the seller trades only when it faces the large liquidity shock;

k) The buyer offers $v_l - c$ to the dealer for bonds in dealer inventory, and buys only low value bonds;

l) When matched, the buyer purchases low value bonds on an agent basis at a price of $v_l - c'$;

m) Total transaction volumes equal $(1 + \theta) * q + 2 * \theta * (1 - q) * p'$

n) Realized bid-offer on round trip principle trades equals $K - c$, and on agent trades equals $\Delta_d - c'$.

Transparency allows the buyer to differentiate between the high and low value bonds, and thus the dealer is able to sell the high value bond via both principle and agent trading. This implies that the dealer bids using the same menu for both types of bonds.

Lemma 4: With high cost, transparency, and high adverse selection, the unique equilibrium is:

f) The dealer offers the seller a choice of $v - K$ with certainty or the ability to "work an order" for $v - \Delta_d$ (which has success rate p'), for K defined using [5] above;

g) The seller chooses immediacy when it faces a large liquidity shock and the order when it faces a low liquidity shock;

h) The buyer buys bonds in dealer inventory at $v - c$;

i) The buyer trades on order (when a match is found) at a price of $v - c'$;

j) Total transaction volumes equal $2 * q + 2 * (1 - q) * p'$

k) Realized bid-offer on round trip principle trades equals $K - c$, and on agent trades equals $\Delta_d - c'$.

Transparency increases volumes, because the dealer is able to sell high value bonds to the buyer, which otherwise either sat in inventory (when the liquidity shock was high) or did not trade (when the liquidity shock was small).

However, the situation is more complex when high costs are combined with low adverse selection. This combination raises an intriguing possibility: that the dealer is willing to offer differential liquidity for high and low value bonds. When adverse selection is low, the buyer is potentially willing to offer.

Empirical Appendix

A1. Figures

Figure A 1: Distribution of bond age before and after the "no publication" event

The figures show the distribution of bond age for transactions reported with transparency.

Panel A: 1st November 2022 – 15th November 2022 (**regular liquidity publication**)

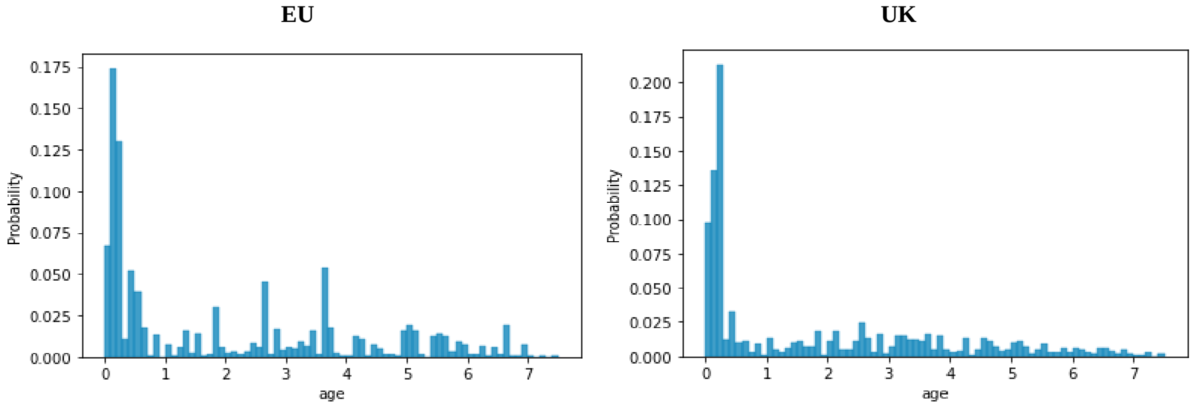

Panel B: 16th November 2022 – 15th February 2023 (**no liquidity publication**)

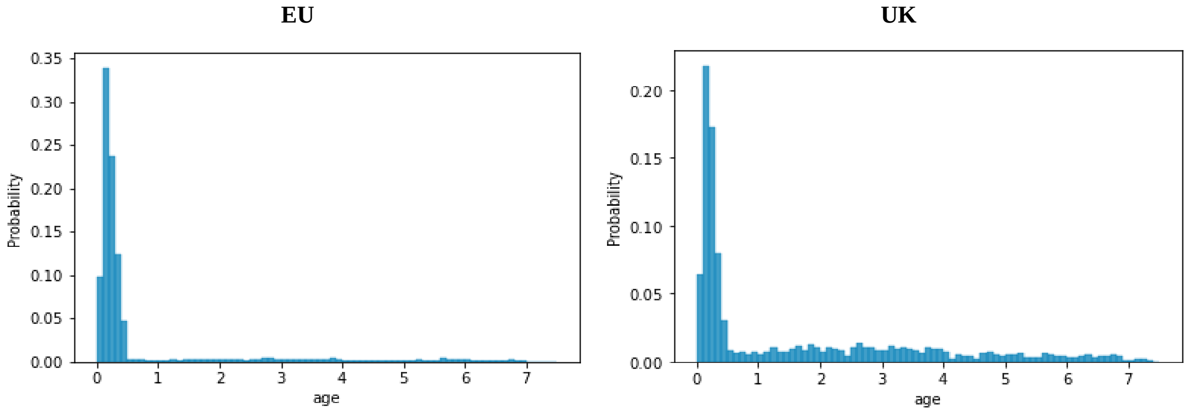

Panel C: 16th February 2023 – 15th May 2023 (**regular liquidity publication**)

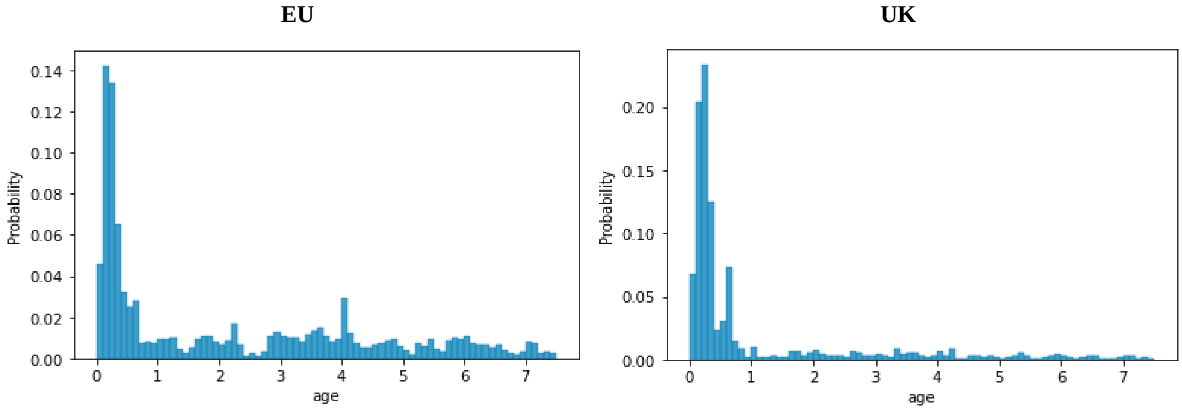

Figure A 2: Distribution of monthly bond turnover, by bond age

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

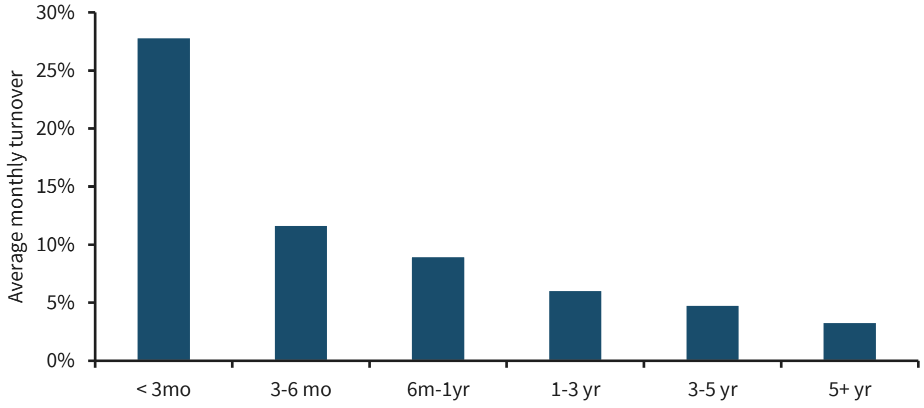

Figure A 3: Distribution of trading activity, by size buckets

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

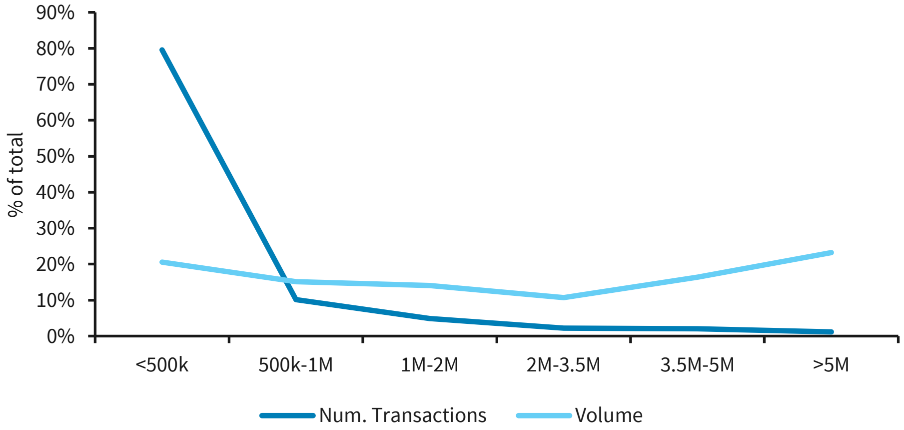

Figure A 4: Parallel trends – agent trading

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

Panel A: Entering the "no publication" period

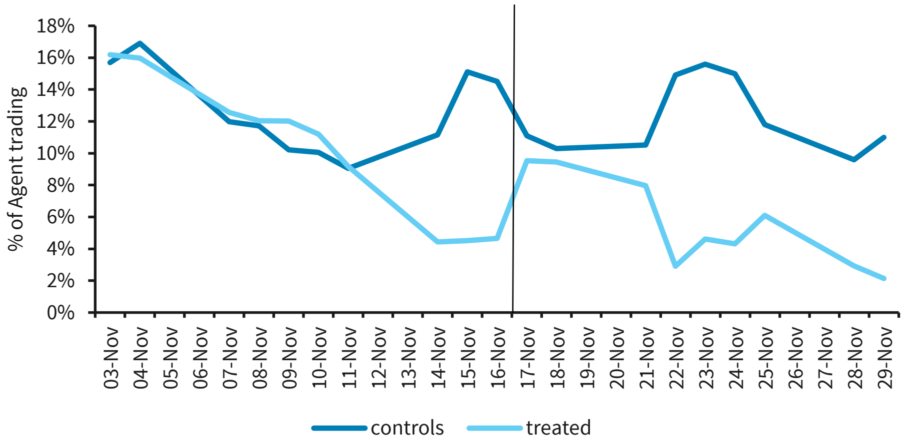

Panel B: Exiting the "no publication" period

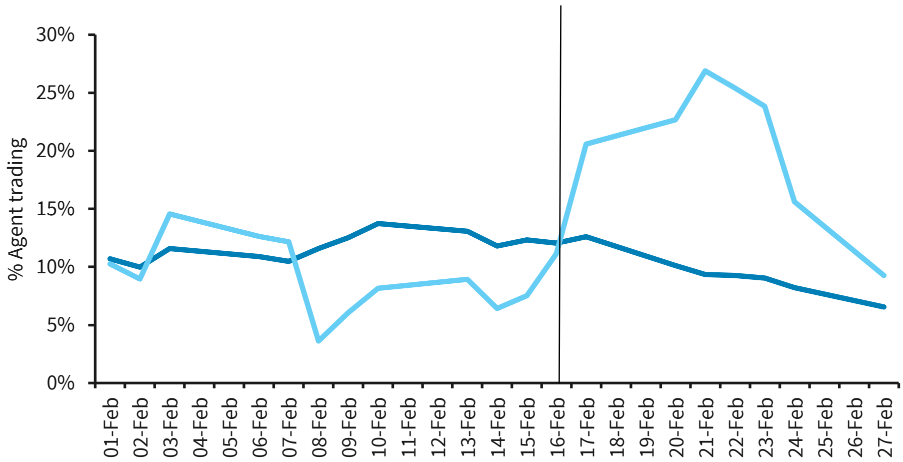

Figure A 5: Parallel trends – bid-offer

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

Panel A: Entering the "no publication" period

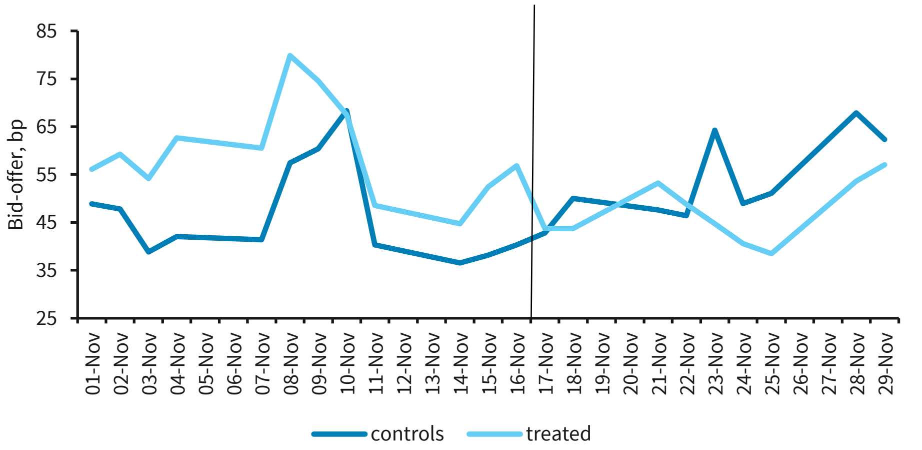

Panel B: Exiting the "no publication" period

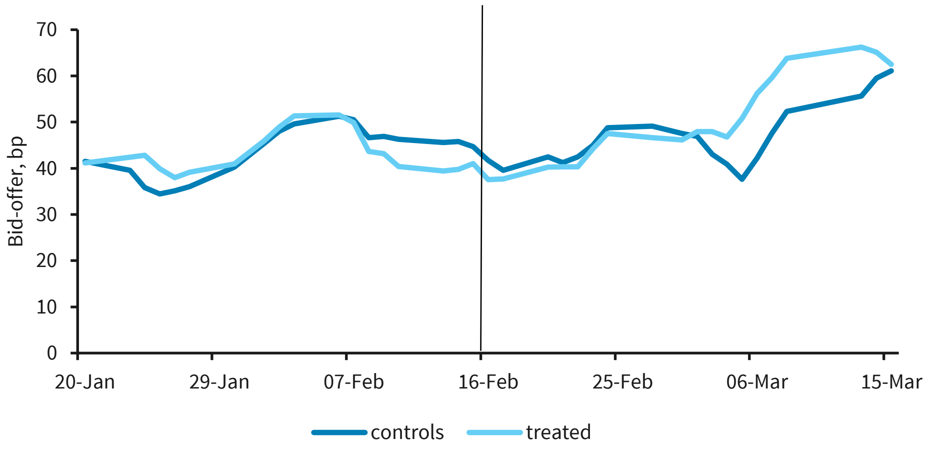

A2. Tables

Table A 1: List of Trading Venues

This table contains a list of the trading venues we used to collect the data

Jurisdiction	Mifid Entities	Mifid Entities Trading Venue	Venue of Publication
EU	AFS	AFS - OTF - BONDS	AFSO
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	BONDSPOT	BondSpot S.A.	TBSP
EU	BONDVISION	MTS S.P.A. - Bond Vision Europe	SSOB
UK	BONDVISION	BONDVISION UK	BVUK
EU	CBOE	CBOE Europe B.V.	CEUX
EU	EBM	EBM - MTF	EBMX
UK	JSE	JSE Limited - Bond Electronic Trading Platform	JSEB
UK	KYTE	Kyte Broking Limited	KBLM
UK	LEDGEREDGE	LEDGEREDGE LTD	LELE
EU	LIQUIDNET	TP ICAP (EUROPE)	LNFI
EU	LIQUIDNET	LIQUIDNET EU LIMITED FIXED INCOME MTF	LEUF
UK	LIQUIDNET	Liquidnet Europe Fixed Income	LIQF
UK	LSE	London Stock Exchange Non-AIM MTF	XLOM
UK	LSE	LONDON STOCK EXCHANGE	XLON
UK	MARIANA	Mariana UFP OTF	MUFP
EU	MARKET SECURITIES	MARKET SECURITIES (FRANCE) SA	MKTF
EU	MARKETAXESS	MarketAxess NL B.V.	MANL
UK	MARKETAXESS	MarketAxess Europe MTF	MAEL
EU	MTS	MTS Belgium	BMTS
EU	MTS	MTS Denmark	MTSD
EU	MTS	MTS Finland	MTSF
EU	MTS	MTS France SAS	FMTS
EU	MTS	MTS S.P.A. - MTS Italia	MTSC

EU	MTS	MTS S.p.A. - MTS Cash Domestic MTF	MCAD
EU	MTS	MTS INTERDEALER SWAPS MARKET	MSWP
UK	SQUARE	Square Global Markets	SQUA
EU	TRAD-X	TRAD-X	TRXE
EU	TRADECHO	UnaVista TRADEcho B.V.	ECEU
UK	TRADECHO	London Stock Exchange plc	ECHO
EU	TRADEWEB	Tradeweb EU B.V.	TWEM
EU	TRADEWEB	Tradeweb EU B.V.	TWEO
UK	TRADEWEB	Tradeweb Europe Limited MTF	TREU
UK	TRADEWEB	Tradeweb OTF	TREO
EU	TRADEWEB APA	Tradeweb EU B.V.	TWEA
UK	TRADEWEB APA	Tradeweb Europe Limited	TREA
EU	TRADITION	TSAF OTC	TSAF
EU	TRADITION	Tradition España OTF	TEUR
UK	TRADITION	Tradition OTF	TCDS
EU	TRAX	MarketAxess Post-Trade B.V.	TRNL
UK	TRAX	Xtrakter Limited	TRAX
EU	TURQUOISE	Turquoise Global Holdings Europe B.V.	TQEA
EU	TURQUOISE	Turquoise Global Holdings Europe B.V.	TQEM
EU	TURQUOISE	Turquoise Global Holdings Europe B.V.	TQEX
UK	TURQUOISE	Turquoise Lit Auctions	TRQA
UK	TURQUOISE	Turquoise Litâ„¢	TRQX
UK	TURQUOISE	Turquoise Platoâ„¢	TRQM

Table A 2: IRC methodology example

Panel A shows a sample of eight transactions with information on the execution date, ISIN, quantity and price. Panel B shows three pairs of agent round-trip trades identified from Panel A. To illustrate how the methodology works, we focus on Trade 1 from Panel A and match it with the closest transaction in the same bond and the same size. In this case, it is matched with Trade 2. This is the first pair of trades showed in Panel B.

Panel A:

Trade	**Execution Date**	**ISIN**	**Quantity**	**Price**
1	29/11/2023	Bond A	2,000,000	100.63
2	29/11/2023	Bond A	2,000,000	104.12
3	29/11/2023	Bond A	1,000,000	102.54
4	29/11/2023	Bond A	1,000,000	104
5	29/11/2023	Bond B	300,000	91.25
6	29/11/2023	Bond B	300,000	95.25
7	29/11/2023	Bond C	100,000	102.4
8	29/11/2023	Bond C	170,000	99.5

Panel B:

1st leg					2nd leg					IRC, bp
Trade	Execution Date	ISIN	Quantity	Price	Trade	Execution Date	ISIN	Quantity	Price	
1	29/11/2023	Bond A	2,000,000	100.63	2	29/11/2023	Bond A	2,000,000	104.12	3.49
3	29/11/2023	Bond A	1,000,000	102.54	4	29/11/2023	Bond A	1,000,000	104	1.46
5	29/11/2023	Bond B	300,000	91.25	6	29/11/2023	Bond B	300,000	95.25	4

Table A 3: Raw Sample Summary Statistics

The table shows summary statistics of the raw sample (ie. before applying the IRC methodology) of bond-trade observations executed in the EU and the UK.

	(3) EU	(4) UK
Panel A: Trading volume		
Mean Trade size	€903.4K	€932.7K
Total volume	€1.27T	€0.90T
Panel B: Bond characteristics		
Mean Outstanding	€924.1 million	€906.9 million
Mean Age	3.5 years	3.3 years
Mean Maturity	4.0 years	4.5 years
Unique issuers	1,025	1,021
Unique ISINs	5,294	5,254
Bond-trade observations	1.63 million	1.24 million
Period	Nov-2022 – Sept-2023	