

Rethinking Transparency – Evidence from a Quasi-Natural Experiment

Jeffrey Meli †
Barclays, US

Zornitsa Todorova
Barclays, UK

Andrea Diaz
Barclays, UK

14th May 2024

Abstract

We develop a model of market making and demonstrate that the effect of transparency on liquidity depends on both the cost of dealer inventory and the degree of adverse selection. When inventory cost is high and adverse selection is low, transparency shifts more trades into the (uncertain) agency protocol and increases the bid-offer of principal trades that are hard to “match”. We test these predictions with a novel database of European corporate bond transactions, exploiting two sources of exogenous variation in transparency. Transparency increases transaction costs and the use of agency trading for large trades and trades in older bonds.

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 make both theoretical and empirical contributions towards a better understanding of the effect of post-trade transparency on the modern corporate bond market. First, we develop a model of market making with an endogenous choice of trade protocol and demonstrate that the effect of post-trade transparency on liquidity depends on the interaction between the cost of dealer inventory and the degree of adverse selection in the market. When the cost of inventory is low and adverse selection is high transparency unlocks two-sided trading in bonds that would otherwise sit in dealer inventory, increasing liquidity and the welfare of both investors and dealers. In contrast, when the cost of inventory is high and adverse selection is low, which we believe describes the current corporate bond market, transparency can reduce liquidity: it pushes more trades into the (uncertain) agency protocol and increases the bid-offer of principal trades that are hard to “match” in agency trading.

Our second contribution is empirical. We construct a novel dataset of transactions in European corporate bonds, which is an understudied market, due to the lack of consolidated TRACE-like reporting. Using this data, we exploit two recent exogenous sources of variation in transparency to test our predictions. We find that transparency both increases the use of agency trading and increases the bid-offer spread of transactions that are difficult to match, such as larger trades and trades in older bonds.

Our results challenge the truism that transparency improves liquidity, which is based on a series of studies examining how the introduction of TRACE affected the US

corporate bond market in the early 2000s. However, much has changed in the corporate bond market in the ensuing decades, and two changes in particular could alter the effect of transparency on liquidity. The first change is reasonably well-understood (and studied in the literature): post-crisis reforms have raised the cost of dealer inventory.¹ The second change, less studied but just as critical in our framework, is a decline in adverse selection in the corporate bond market due to the rapid growth of corporate bond ETFs. Before these instruments existed, trade reporting was the only reliable source of information about bond valuation. Now ETFs provide investors and market makers real-time clarity about the price of diversified baskets of corporate credit risk. This reduces uncertainty about valuation to security-specific factors.²

To study the interaction between these forces, we build a model of market making with an endogenous choice of trade protocols. Our framework is loosely based on Saar et al. (2022), who study the effect of an increase in the cost of inventory in a transparent market. Like with that study, in our model raising the cost of inventory increases both the cost of principal trades and the use of agency trading. We demonstrate that the magnitude of these changes depends on the probability of a “match” in the agency protocol. Securities that are easy to match are more insulated from the negative effects of a higher inventory cost, both because agency trading is itself more effective, and because the

¹ The figure is based on net positions in corporate bonds for US primary dealers, available through the Federal Reserve Bank of New York [Primary Dealer Statistics](#) database. Although similar data is not readily available in Europe, we expect a similar pattern exists for European market-makers.

² This can include sector, rating, and tenor factors, as well as security-specific factors. That said, in some markets, like the US, narrow ETFs reduce adverse selection even further. Of course, security-level uncertainty about valuation persists; our assertion is that these instruments reduce adverse selection vis-à-vis the period when TRACE was introduced.

availability of effective agency trading limits how much dealers can increase the cost of principal trades in response to a rise in the cost of inventory.

We then introduce uncertainty in security valuation, which allows us to examine the effect of transparency on liquidity in different regimes. When the cost of inventory is low and adverse selection is high, adverse selection inhibits some two-sided trading.

Inventory incurs only a small cost, and uncertainty about security valuation significantly depresses the price less-informed investors are willing to pay for securities in dealer inventory. Therefore, dealers prefer to hold some positions rather than sell them.

Transparency unlocks two-sided trading, improving the welfare of investors and dealers.

The dynamic changes when a high inventory cost is combined with low adverse selection. In this setting a lack of transparency no longer (necessarily) inhibits two-sided trading. Dealers have a greater incentive to avoid incurring the cost of inventory, and less-informed investors are willing to pay a higher price for securities in dealer inventory. In fact, this second effect can encourage liquidity provision; dealers can make incremental profits in principal trades in some securities. This offsets the drag on liquidity imposed by high inventory costs, such that dealers execute more principal trades with a small bid-offer, and fewer agency trades, as they did in the low cost equilibrium. Introducing trade reporting undoes this effect, forcing dealers to internalize their full inventory cost. The result is higher bid-offer spreads on principal trades and greater use of the agency protocol. Importantly, this effect is not uniform across trades. The reduction in liquidity associated with transparency is more severe for trades that are difficult to match in the agency protocol, which are the trades that are most affected by an

increase in the cost of inventory. Easy to match trades are insulated from the effects of a higher inventory cost, and thus are less sensitive to transparency.

The low cost, high adverse selection and the high cost, low adverse selection regimes characterize the pre-GFC and post-reform (and current) markets. This leads to our main hypothesis: under currently prevailing market conditions, transparency increases bid-offer spreads on principal trades and shifts more transactions into the agency protocol. Further, the magnitude of these effects will be larger for transactions that are difficult to match in the agency protocol, such as larger trades and trades in older bonds.

To test our hypotheses, we need a modern TRACE-like transparency experiment. We turn to Europe, where the MiFiDII reforms that took effect in January 2018 provide a unique setting to study the effects of transparency. MiFiDII introduced trade reporting for a wide set of asset classes, including for corporate bonds. While no consolidated tape exists, the trades are public, and we assemble the first (so far as we are aware) comprehensive set of dealer-to-client trades in European corporate bonds, accumulated from a large number of different voice and electronic venues. We use this dataset to estimate the bid-offer spread of round-trip transactions, distinguishing between principal and agency trades.

We exploit two sources of exogenous variation in the transparency of corporate bond trades executed in the EU and in the UK: a data issue that changed which trades were subject to real-time reporting in the EU (but not UK) for one quarter, and Brexit, which caused a divergence in the reporting obligations for trades in the same bonds based on

jurisdiction.³ These allow us to isolate the effect of transparency by comparing trades in the same bonds but with different reporting obligations. The results align with our predictions. First, we use the one-quarter change in EU reporting in 2022 to compare “treated” bonds (those with a change in transparency) and “control” bonds (those with no change) using two sets of difference-in-difference regressions: we assess the impact of the initial change in reporting, and the effect of restoring the original rules. We demonstrate that the proportion of agency trades in treated bonds declined (*vis-à-vis* control bonds) when transparency was disrupted, and then increased once transparency was restored. Similarly, the bid-offer of principal trades in treated bonds decreased when transparency was disrupted and increased once it was restored. Second, we use the variation linked to Brexit to demonstrate that the effect of transparency on the bid-offer of principal trades depends on the size of the position and the age of the bond; for smaller trades and younger bonds, transparency slightly reduces bid-offer (e.g., by 5% for small trades), but for larger trades and trades in older bonds, transparency increases bid-offer (e.g., by 15% for large trades).

Our analysis has important policy implications. Europe is revising its trade reporting requirements, and our results suggest that a fully transparent market does not lead to the highest liquidity in the current environment. The need for “radical” transparency has declined as other sources of information have become available to investors. Instead, some amount of opacity can increase intermediation, particularly for trades that are hard to match. These trades are costly for dealers because they are likely to spend some time

³ Investors trade where domiciled; they have no discretion as to the jurisdiction they are subject to.

in inventory. Some version of the current sliding scale, where smaller trades and trades in more liquid securities are reported in real time, but larger trades, and trades in older, less liquid bonds are reported with a delay, is optimal. In addition, our analysis implies that the optimal transparency regime will vary across markets, depending on both the cost of inventory associated with those securities and the degree of adverse selection. For example, there is concern that an increased inventory cost (linked to the leverage ratio) has reduced liquidity in the US Treasury market; regulators are considering a proposal to extend post-trade reporting to off-the-run Treasuries as a possible mitigant. Given that the degree of adverse selection in that market is likely low, increased transparency may do more harm than good. The situation may be reversed in an opaque asset class, such as leveraged loans. These are interesting avenues for future work.

Literature Review

We contribute to two strands of the literature. First, our work relates to theoretical models of transaction costs, liquidity, and transparency in over-the-counter (OTC) markets, including Duffie, Dworczak, & Zhu (2017), and Vairo & Dworczak (2023).⁴ The conclusions are mixed. Duffie et al. (2017) find that transparency (in the form of a published benchmark) typically increases liquidity. Vairo & Dworczak, (2023) conclude that pre-trade transparency leads to more efficient outcomes than post-trade transparency. Back, Liu, & Tegua, (2018) assume the market maker faces an infinite cost of inventory, and so never holds any positions; the authors find benefits of transparency. However,

⁴ See also Duffie, Gârleanu, & Pedersen, (2005), who consider the effect of search and bargaining on valuation in OTC markets, and Zhu (2012), who studies the effect of a repeat contact with the same buyer, and Glebkin, Yueshen, & Shen, (2022).

none of these models includes an ability to transact via both agency (or “matchmaking”) and principal trading, which is increasingly common in the corporate bond market (Goldstein & Hotchkiss, (2020), Choi, Huh, & Shin, (2023)).

Two recent models that do incorporate agency and principal trading are An & Zheng (2022) and Saar et al. (2022), both of which analyze transparent markets. The former examines how dealers prioritize positions in inventory versus matching customer orders. The latter, which is the closest to our model, examines how the effect of inventory costs on investor welfare depends on the competitive dynamics faced by dealers. Our model is most similar to their “unconstrained bank” scenario, and we also find that a higher inventory cost leads to more agency trading and more costly principal trades. We introduce uncertainty about security valuation, allowing us to analyze how transparency affects liquidity in different market environments.

We also contribute to the empirical literature that studies the effect of transparency on 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., Bessembinder, Maxwell, & Venkataraman, (2006); Goldstein, Hotchkiss, & Sirri, (2006); Edwards, Harris, & Piwowar, (2007); Bessembinder & Maxwell, (2008); Asquith, Covert, & Pathak, (2019)). The general conclusion of these papers is that transparency decreases transaction costs for investors and increases trading volumes. Edwards, Harris, & Piwowar, (2007) show that while transaction costs decline more significantly for small trades due to transparency, there is still a noticeable reduction in transaction costs across all bond types. This finding is consistent with our theoretical model, which posits that

transparency decreases transaction costs when inventory costs are low and adverse selection is high, as is the case when TRACE was introduced.

Our data and experiment are novel for three reasons. First, it is from the European corporate bond market, which remains understudied due to the lack of consolidated trade reporting. The construction of a clean consolidated tape in European corporate bonds from various reporting venues is a significant contribution in itself. Second, our data is current, spanning the period from November 2022 to September 2023, providing fresh empirical evidence on how post-trade transparency impacts transaction costs under the now-prevalent conditions in the corporate bond market. Third, the design of our experiment enables us to establish a more robust pre-transparency sample for “treated” and “control” bonds, as reporting exists even in the absence of transparency. In our context, transparency means that reporting is delayed, but all trades are eventually observed. This contrasts with previous studies utilising the introduction of TRACE, where no reporting existed prior TRACE, requiring constructing the “pre” sample from proprietary data, which was typically limited to trades from a single investor or a smaller number of bonds.

2. Model

2.1. Primitives

Players: There are three risk neutral players, which each maximize their expected payoff.

“Seller”: An investor that owns a security of value v equal to either v_l or v_h with probabilities θ and $(1 - \theta)$ respectively, and $v_l < v_h$. The seller experiences a liquidity

shock Δ , such that the private value of the security to the seller is $v - \Delta$. The liquidity shock is Δ_u or Δ_d , with probabilities q and $(1 - q)$ respectively, with $\Delta_u > \Delta_d > q * \Delta_u$. The last inequality ensures that the dealer would satisfy all liquidity needs absent frictions.

“Dealer”: A market maker in the security, willing to provide liquidity to the seller. The dealer can either sell the security once acquired (see below), or hold it in inventory, at a cost $c \geq 0$. As in prior studies, we assume the dealer has market power and thus makes take-it-or-leave-it quotes to potential counterparties.

“Buyer”: An investor who is potentially willing to buy the security at a price negotiated with the dealer. A buyer arrives with probability p . The buyer also experiences a liquidity shock Δ (also equal to either Δ_u and Δ_d , with probabilities q and $1 - q$), such that the value of the security to the buyer is $v + \Delta$.

The liquidity shocks are a source of potential gains from trade. We assume that the shocks experienced by the seller and buyer are uncorrelated with each other and with security type, and are unknown to any other players (e.g., only the seller knows its liquidity shock). The value of the security is known to both the seller and the dealer. The information available to the buyer depends on the transparency regime. Absent transparency, the buyer knows only the distribution of values and the cost of inventory. However, when we impose transparency via trade reporting, the buyer can infer the value of the security from transaction reporting. Finally, we assume that all players have a weak preference for trading.

Timing

There is one period, with the following stages. First, the seller asks the dealer for a bid on the security; the dealer responds with a take-it-or-leave-it quote (see below for the various forms the quote can take). The bid is a function of the security value, which is known to the dealer. If the dealer purchases the security or accepts an order, it will make a take-it-or-leave it ask to the buyer, if one arrives. We define the bid-ask spread to be the average of the difference between the purchase and sale prices for round-trip trades; this may vary by trade protocol. Finally, at the end of the period all players realize their payoffs, which for the dealer includes a cost incurred for any inventory it holds.

Trading and equilibrium

Two trading protocols are available to the dealer. We label the first protocol “principal trading”; the dealer trades sequentially with the seller and the buyer, and holds any unsold positions in its inventory.

Second, we consider an “agency trading” protocol, in which the dealer attempts to pre-arrange both the sides of the trade and execute them simultaneously. Agency trading is only successful if a buyer arrives, which occurs with probability p .⁵ Therefore, agency trading requires the seller to sacrifice certainty.⁶ We assume that the dealer cannot “lie”

⁵ Note that the buyer arrival process is unrelated to the trade protocol. Dealers have one salesforce (as opposed to one salesforce that exerts effort to find matches and another that handles principal trades) and one set of clients, all of whom are willing to trade in either protocol.

⁶ Agency trading has two frictions vis-à-vis principal trading: timing and certainty. Other models, such as Saar et. al. (2022), focus on the timing of execution; we use a one-period model and so the risk of agency trading is that some orders are unfilled.

to the buyer about the terms of the transaction, which implies that in an agency trade the buyer learns the value of the security, even absent transparency.⁷

The dealer chooses to make a principal-only bid, an agency-only bid, or provide a “menu” of both protocols (at different prices). Due to the possibility of unmatched agency orders, a menu will involve a trade-off: certain, but expensive, principal trading, versus uncertain, but cheap, agency trading. The dealer selects the trade protocol(s) and price(s) to quote the seller and buyer that maximize its expected profits, conditional on the buyer and seller decision rules, as detailed below.

Equilibrium is defined as a set of strategies for the three players that constitutes a “trembling hand perfect” Nash equilibrium.

Seller and buyer decision rules

Both the seller and the buyer maximize their expected payoff when choosing if and how to transact with the dealer. Because the seller knows the value of the security, it trades so long as the highest bid price it receives from the dealer is at least equal to its private value:

$$\max(B_p, B_a) \geq v - \Delta \rightarrow trade \quad [1]$$

⁷ We make this assumption for analytical convenience; our conclusions are unchanged if the buyer does not learn the value of the security in an agency trade. In the context of our data, our assumption is more appropriate, because the European regime does eventually release trade details; reduced transparency only delays this revelation. Therefore, given the market standard that the dealer states its purchase price in an agency trade, it is unlikely that a dealer could sustainably prevaricate.

where B_p and B_a are the principal and agency bids it receives from the dealer (equal to 0 if the dealer does not offer that protocol). It chooses the principal trade if it generates a greater improvement in expected utility (accounting for the uncertain agency execution):

$$B_p - (v - \Delta) \geq p * [B_a - (v - \Delta)] \rightarrow \text{choose principal trade} \quad [2]$$

The inequalities in [1] and [2] reflect the weak preference to trade.

The buyer does not (necessarily) know the value of the security. Its decision to trade is linked solely to the ask price A :

$$A \leq E^*(v) + \Delta \rightarrow \text{trade} \quad [3]$$

The expectation in [3] is conditional: it depends on both the trade protocol and the equilibrium strategies. For example, we assume that in an agency trade the dealer (credibly) discloses its purchase price, which reveals the value of the security, implying that $E^*(v) = v$. The same holds for all trades under transparency. For principal trades done without transparency, however, this expectation will reflect the distribution of securities the buyer is offered from dealer inventory, which may differ from the unconditional distribution of security type.

The combination of trembling hand perfect equilibria and the weak preference for trading ensure that these are the only viable equilibrium decision rules. They maximize seller and buyer utility and the probability of trading. Thus, they (at least) weakly dominate other decision rules that could comprise a Nash equilibrium. For example, we can construct a Nash equilibrium where the dealer always bids a price that reflects the large liquidity shock, by pairing it with a decision rule where the seller only trades when

it experiences the large liquidity shock. Neither player has an incentive to deviate (e.g., the seller has no incentive to deviate because the dealer never offers a price that reflects the low liquidity shock). However, the seller strategy does not survive a “tremble” bid that reflects the low liquidity premium; transacting at that price either increases utility or leads to a zero-utility trade (depending on the seller’s liquidity shock). Therefore, both the seller and buyer must be willing to transact at any price that (at least weakly) increases utility. This insight simplifies solving for the equilibria: we need only determine the optimal strategy for the dealer given these decision rules.

Lemma 1: The above decision rules are the only viable equilibrium strategies.

2.2. Transparency

We first solve for the equilibria that obtain in a transparent market. We consider “low” and “high” cost regimes. These will establish a baseline for comparison once we remove transparency and confirm that the implication of an increase in the cost of inventory in our model comports with the conclusions of prior studies.

Transparency allows the buyer to infer the value of the security by observing the price of the transaction between the seller and the dealer; thus, we drop the subscripts from the security value v . It also constrains the quotes the dealer will make in equilibrium. The ask is the most constrained. Due to the discrete nature of the liquidity shock, the only sustainable asks are $v + \Delta_d$ and $v + \Delta_u$. The buyer will transact at any price $A \leq v + \Delta_d$, regardless of its liquidity shock, but in this range the dealer’s profits are greatest at $v + \Delta_d$. Similarly, the buyer will only purchase the security at a price $v + \Delta_d < A \leq$

$v + \Delta_u$ if $\Delta = \Delta_u$, but in this range dealer profits are greatest at $v + \Delta_u$. The buyer does not transact at any price $A > v + \Delta_u$.

Second, we demonstrate that, between these two choices, the dealer always prefers to quote $A = v + \Delta_d$ (the dealer only provides an ask if the buyer arrives). Once the dealer has executed a principal trade with the seller, it has the security in its inventory. A buyer will always purchase the security at $v + \Delta_d$. A quote of $v + \Delta_u$ will trade with probability q ; otherwise, the dealer must hold the security in inventory. The dealer's profit is always greater in the former case:

$$\begin{aligned} \Delta_d &> q * \Delta_u - (1 - q) * c \rightarrow \\ (\Delta_d - q * \Delta_u) + (1 - q) * c &> 0 \end{aligned} \quad [4]$$

Because the inventory cost is nonnegative and $\Delta_d > q * \Delta_u$, [4] is always satisfied.

Similarly, if the dealer has agreed to an agency trade with the seller at a discount of $X \geq 0$ to intrinsic value, the dealer's profits are higher if it quotes the buyer $v + \Delta_d$:

$$\begin{aligned} \Delta_d + X &> q * (\Delta_u + X) \rightarrow \\ (\Delta_d - q * \Delta_u) + (1 - q) * X &> 0 \end{aligned} \quad [5]$$

Although there is more flexibility in the bids, each of the three types (agency-only, principal-only, and a menu) are constrained. An agency-only bid must be $v - \Delta_d$. To see why, compare the dealer's profits from $v - \Delta_d$ and $v - \Delta_u$ in an agency-only bid, knowing that the buyer will purchase the security at $v + \Delta_d$ if it arrives:

$$2p * \Delta_d > p * q(\Delta_u + \Delta_d) \rightarrow$$

$$\Delta_d * (1 - q) + (\Delta_d - q * \Delta_u) > 0 \quad [6]$$

Once again, [6] is true by assumption. By extension, the expected profitability of an agency-only bid is $2p * \Delta_d$.

Similarly to the asks, a principal-only bid must be either $v - \Delta_d$ or $v - \Delta_u$ (which trades with probability q). We compare the profitability of these bids, knowing that the security will be sold to a buyer (if it arrives) at $v + \Delta_d$. The expected profit of a principal bid of $v - \Delta_d$ is:

$$\pi(v - \Delta_d) = \Delta_d * (1 + p) - c * (1 - p) \quad [7]$$

where [7] reflects both the probability of selling the security to the buyer (and conversely of holding it in inventory) and the associated profits of the round-trip trade. The expected profit of a principal bid of $v - \Delta_u$ is:

$$\pi(v - \Delta_u) = q * [\Delta_u + p * \Delta_d - c * (1 - p)] \quad [8]$$

Comparing [7] and [8], we find an inventory cost threshold below which the bid of $v - \Delta_d$ is more profitable:

$$\begin{aligned} \pi(v - \Delta_d) &\geq \pi(v - \Delta_u) \rightarrow \\ \Delta_d * (1 + p) - c * (1 - p) &\geq q * [\Delta_u + p * \Delta_d - c * (1 - p)] \rightarrow \\ c &\leq \frac{(\Delta_d - q * \Delta_u)}{[(1-p) * (1-q)]} + \Delta_d * p / (1 - p) \end{aligned} \quad [9]$$

According to [9], at a high enough cost, the consequence of holding the bond in inventory is too large to risk trading unless the dealer increases the transaction cost it

charges the seller. It will still quote $v + \Delta_d$ to the buyer, but the dealer prefers to do fewer trades at a higher profitability to offset the larger (expected) drag from inventory.

Finally, turning to the menu of principal and agency trading, it takes the form:

$$\begin{aligned}
 B_p(v) &= v - K \text{ with certainty or} \\
 B_a(v) &= v - X \text{ if a match is found}
 \end{aligned}
 \tag{10}$$

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

$$\begin{aligned}
 \Delta_u - K &= p(\Delta_u - \Delta_d) \rightarrow \\
 K &= (1 - p)\Delta_u + p\Delta_d
 \end{aligned}
 \tag{11}$$

The important insight in [11] is that the required discount to the cost of principal trading increases with the probability of a match in the agency market. For example, if $p = 1$, $K = \Delta_d$, meaning that the dealer prices principal trades equivalently to agency trades.

Low inventory cost

We first assume that the inventory cost meets the criteria in [9], such that the optimal principal-only bid is $v - \Delta_d$. We compare the profitability of that strategy to an agency-only bid (also of $v - \Delta_d$), which will trade with a probability of p but has no risk of incurring an inventory cost. The principal-only bid is preferred when:

$$\begin{aligned} \Delta_d * (1 + p) - c * (1 - p) &\geq 2p * \Delta_d \rightarrow \\ c &\leq \Delta_d \end{aligned} \tag{12}$$

Next, we compare the principal-only bid to the menu of agency and principal bids designed to induce separation. The principal-only bid is preferred when:

$$\begin{aligned} \Delta_d(1 + p) - c * (1 - p) &\geq 2p(1 - q)\Delta_d + q * [K + p * \Delta_d - c * (1 - p)] \rightarrow \\ c &\leq (\Delta_d - q * \Delta_u)/(1 - q) = c' \end{aligned} \tag{13}$$

where we substitute for K as defined in [11].

The constraint in [13] is tighter than that in [9] (because the menu is preferred over a principal-only strategy above c'). We define a “low” inventory cost to be below these thresholds.

Definition: The cost of inventory is “low” when $c \leq \min(\Delta_d, c')$.

A low cost implies that the optimal strategy for the dealer is a principal-only bid of $v - \Delta_d$, allowing us to describe the resulting equilibrium.

Lemma 2: The unique low cost, transparent equilibrium is:

- a) *The dealer offers $v - \Delta_d$ to the seller, and buys all securities on a principal basis;*
- b) *The buyer (if it arrives) buys securities at $v + \Delta_d$ from the dealer;*
- c) *Total transaction volume equals $1 + p$;*
- d) *Realized bid-offer on round trip trades equals $2 * \Delta_d$;*
- e) *Dealer inventory is $1 - p$.*

Trading volumes equal $1 + p$ because all securities are bought by the dealer on a principal basis and are sold to the buyer, if it arrives. This is the highest possible volume; all realized liquidity shocks are met. Similarly, the bid-offer of $2 * \Delta_d$ is the lowest possible transaction cost.

This is a high liquidity equilibrium. The dealer facilitates all possible trades via principal trading at a low bid-offer, and warehouses inventory of unsold positions. The expected welfare of both the seller and the buyer are $q * (\Delta_u - \Delta_d)$ because they execute at a low bid-offer even when they experience the large liquidity shock. Qualitatively, this equilibrium corresponds to the pre-GFC status quo: trading volume was high, bid-offer was low, dealer inventory was high, and principal trading dominated.

High inventory cost

We now consider how the equilibrium changes when the inventory cost is high. To start, we assume that $c > c'$, such that the menu is more profitable than the principal-only bid of $v - \Delta_d$. For the menu to be the equilibrium strategy, it must generate greater profits than both a principal-only bid of $v - \Delta_u$ (which trades with probability q) and the agency-only bid. The former is always true:

$$2(1 - q) * p * \Delta_d + q[K + p * \Delta_d - c * (1 - p)] \geq q * [\Delta_u + p * \Delta_d - c * (1 - p)] \rightarrow$$

$$(\Delta_d - q * \Delta_u) + (1 - q) * \Delta_d \geq 0 \quad [14]$$

The latter imposes an upper bound on the inventory cost:

$$2(1 - q) * p * \Delta_d + q[K + p * \Delta_d - c * (1 - p)] \geq 2 * p * \Delta_d \rightarrow$$

$$c \leq \Delta_u \quad [15]$$

The constraint in [15] ensures that the inventory cost is not excessively large such as to deter all principal trading.⁸

Definition: Inventory cost is “high” if $c \in (c', \Delta_u]$.

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

- a) *The dealer offers the seller a choice of a certain principal trade at $v - K$ or an agency trade at $v - \Delta_d$ (which has success rate p), for K defined in [11];*
- b) *The seller chooses immediacy when it faces a large liquidity shock and the agency protocol when it faces a small liquidity shock;*

⁸ Of course, principal trading continued to occur even after the post-crisis reforms.

- c) The buyer purchases the security at $v + \Delta_d$ if it arrives;
- d) Total transaction volumes equal $q + p * (2 - q)$;
- e) Dealer inventory equals $q * (1 - p)$;
- f) Realized bid-offer on round trip principal trades equals $K + \Delta_d$;
- g) Realized bid-offer on round trip agency trades equals $2 * \Delta_d$

By comparing the equilibria in *Lemma 1* and *Lemma 2*, we see that the effect of an increase in inventory cost comports with both the post-GFC experience and the conclusions of similar models.⁹ First, a higher cost leads to an increase in agency trading. Agency trading plays a specific economic function: it facilitates trades that would otherwise be precluded by the high cost of inventory.

Second, the volume in *Lemma 2* reflects the choice of protocol: the dealer purchases securities on a principal basis when the seller experiences a large liquidity shock and sells them to the buyer when it arrives. It also matches both sides of an agency trade with probability p when the seller experiences the small liquidity shock. The use of agency trading reduces trading volume vis-à-vis *Lemma 1*, as some agency orders are unfilled. It also reduces dealer inventory, which only arises when principal bids are unmatched.

The bid-offer on round-trip trades in *Lemma 2* varies by protocol. The bid-offer on principal trades is higher than in *Lemma 1* because the dealer bid (partially) reflects the

⁹ Most notably, this corresponds to the effects of an increase in inventory costs in the unconstrained bank scenario in Saar et. al. (2022).

large liquidity shock. That said, agency trades have a relatively low bid-offer, which ameliorates the effect of inventory costs on the average bid-offer of all round-trip trades.

The increased cost of inventory reduces seller welfare for two reasons: some agency orders are unfilled, and the transaction cost it pays on principal trades increases. Dealer profits decline because it does fewer trades and has a higher expected inventory cost for each principal trade. The welfare of the buyer is unchanged (it buys the security at $v + \Delta_d$ if it arrives in both equilibria).

This aligns closely with the post-crisis corporate bond market. A rise in the cost of inventory reduces trading volume and increases the prominence of agency trading. Average bid-offer increased, but the use of agency trading mitigates the observed effect on transaction costs. Unfilled (but unobserved) agency orders reduce investor welfare.

Importantly, the magnitude of these effects is driven by the probability of finding an agency match. The existence of agency trading as an option requires that principal trades be done at a low enough bid-offer to induce separation. This discount on principal trades grows as the probability of a match increases (and thus as K declines) because a higher probability of a match increases the seller's expected utility from the agency protocol when it has the large liquidity shock. In fact, all of the differences between the high and low cost equilibria shrink as the probability of a match increases; put differently, transactions that are easy to match are more insulated against the negative effects of a higher inventory cost on liquidity.

2.3. Removing Transparency: The Implications of Adverse Selection

We will demonstrate how transparency effects liquidity, measured by trading volume and transaction costs. To do so, we introduce uncertainty about security valuation. This affects the buyer, which can no longer (in a principal trade) differentiate between the two types of securities. It must base its decision on the expected value of securities it is quoted from the dealer's inventory. In equilibrium, this expectation must be correct: the dealer must behave exactly as the buyer predicts (and vice-versa).

There are two types of equilibria. The first is a separating equilibrium, where the outcome varies by security type. The second is a pooling equilibrium, where the buyer purchases both types of securities at the same price. Our first step is to find a sufficient condition for ruling out pooling equilibria and show how it is linked to both the uncertainty about security valuation and the cost of inventory.

Recall that the buyer decision rule is:

$$A \leq E^*(v) + \Delta \rightarrow \text{trade} \quad [16]$$

The expectation in [16] is conditional on the distribution of security types that the dealer will sell out of inventory. As we will see below, this may differ from the population distribution: the low value security may be overrepresented but not underrepresented. Further, the lower the pooled price, the less likely a pooling equilibrium, because the dealer must prefer selling the high value security at the pooled price, rather than holding it in inventory. Therefore, in furtherance of finding a sufficient condition for ruling out pooling, we assume (for now) that the conditional expectation in [16] is equal to the

unconditional expectation of security value, which results in the highest possible pooled prices: $E^*(v) = E(v) = \theta * v_l + (1 - \theta) * v_h$.

As above, the dealer chooses the price that maximizes its profits, which implies it chooses between two possible pooling prices: $E(v) + \Delta_d$ and $E(v) + \Delta_u$, the latter of which trades with probability q .

A pooling equilibrium requires that the dealer sell both securities at the same price. This means it prefers selling over holding the security in inventory, and that it will choose the same ask for both types of securities from the two options above.

The dealer always prefers selling the low value security at the pooled price: it is above intrinsic value and selling avoids incurring the cost of inventory. Further, the dealer always prefers to sell the low value security at $E(v) + \Delta_d$:

$$E(v) + \Delta_d \geq q * (E(v) + \Delta_u) + (1 - q) * (v_l - c) \rightarrow$$

$$(E(v) - v_l) + c + c' \geq 0 \quad [17]$$

We turn to the high value security, and test if (i) the dealer prefers to sell it at the pooled price of $E(v) + \Delta_d$ rather than hold it in inventory and (ii) if it prefers to sell it at $E(v) + \Delta_d$ or at $E(v) + \Delta_u$ (with probability q). Each imposes a constraint on the cost of inventory. First, we compare selling at the pooled price to holding it in inventory. The dealer sells at the pooled price if:¹⁰

$$E(v) + \Delta_d \geq v_h - c \rightarrow$$

¹⁰ This is also sufficient to ensure that the dealer does not sell the high value security at the low price.

$$c \geq \theta * (v_h - v_l) - \Delta_d \quad [18]$$

Second, we compare the two possible pooled prices. The dealer prefers the low (but certain) pooled price if:

$$E(v) + \Delta_d \geq q * [E(v) + \Delta_u] + (1 - q) * (v_h - c) \rightarrow$$

$$c \geq \theta * (v_h - v_l) - c' \quad [19]$$

The $\theta * (v_h - v_l)$ term in [18] and [19] is the discount to intrinsic value that the dealer receives for the high value security. The second terms are the incremental gains from trade associated with selling at the lower, but certain, price (incremental vis-à-vis either holding the security in inventory or selling only if the buyer experiences the large liquidity shock). When the inventory cost is below the difference between these the dealer will not sell the high value security at the pooled price. The buyer, knowing this, will not pay the pooled price. Instead, it will (correctly) presume that any security it is quoted at that price has the low value, and thus will pay at most $v_l + \Delta_d$ (if it has the small liquidity shock) or $v_l + \Delta_u$ (if it has the large liquidity shock).

Let $c'' = \max(\theta * (v_h - v_l) - c', \theta * (v_h - v_l) - \Delta_d)$. A cost below this threshold is sufficient for ruling out pooling equilibria.

Corollary: Pooling equilibria are not viable if $c < c''$.

Definition: When $c < c''$ adverse selection is high, otherwise it is low.

This criterion is more likely to be met in the low cost regime, where the consequence of owning the security is limited. In other words, any level of uncertainty

about security valuation can meet this definition if the cost of inventory is low enough.

This further supports our pairing of a low inventory cost with high adverse selection, and vice-versa.

Low cost, high adverse selection

Only separating equilibria are viable in this setting. The buyer will utilize one of two possible decision rules. In the first, the buyer behaves as if any security it is quoted by the dealer is worth v_l :

$$A \leq v_l + \Delta \rightarrow \text{trade} \quad [20]$$

In the second, the buyer is willing to purchase the high value security when it has the large liquidity shock. This decision rule is dependent on the buyer's liquidity shock:

$$\text{for } \Delta = \Delta_d, A \leq v_l + \Delta_d \rightarrow \text{trade}$$

$$\text{for } \Delta = \Delta_u, A \leq v_h + \Delta_u \rightarrow \text{trade} \quad [21]$$

In both, the buyer will pay a price associated with the low value security when it experiences the small liquidity shock because adverse selection is high: the dealer won't sell the high value security at a price the buyer is willing to pay. The difference between the two rules is that the buyer may, under certain conditions, be willing to pay a price associated with the high value security when it experiences the large liquidity shock. This is viable only if the dealer will not "cheat" and sell the low value security at that price. In other words, the dealer must prefer the certainty of selling the low value security at $v_l + \Delta_d$ over selling it at $v_h + \Delta_u$ with probability q . If [21] is viable, the buyer must use it

(else it is forgoing possible trades); we detail the parameter values for which this holds in the Appendix.

Regardless, we can describe the resulting equilibrium (see the Appendix for a formal treatment). Using either [20] or [21], the equilibrium for the low value security follows *Lemma 2*: the dealer purchases it on a principal basis at $v_l - \Delta_d$ and sells it to the buyer (if it arrives) at $v_l + \Delta_d$. The dealer has three options for the high value security. First, it can engage in agency-only trading, which allows it to credibly disclose the security type. Second, it can use principal-only trading, and either hold the high value security in inventory (when the buyer uses [20]) or sell it at $v_h + \Delta_u$ when the buyer experiences the large liquidity shock (when the buyer uses [21]). Finally, it can use the menu, which is a combination of these. Each of these options can be optimal depending on the inventory cost (in the Appendix we detail the relevant cost thresholds). If the buyer uses [21], both the principal-only and the menu are utilized over a wider range of inventory cost because trading the high value security is more profitable.

Any of these options results in lower volume than in the transparent market. At the very least, two-sided principal trading in the high value security does not occur when the buyer experiences the small liquidity shock. This is replaced by either (uncertain) agency trading or one-sided principal trading, where the security remains in dealer inventory. The decline in volumes is more severe when the buyer uses [20] because there are no round-trip principal trades in the high value security.

When the buyer uses [20], the bid-offer of round trip principal trades is $2 * \Delta_d$, the same as in *Lemma 2*. However, if the buyer uses [21], the average round-trip bid-offer of principal trades increases, because the bid-offer of trades in the high value security is $\Delta_d + \Delta_u$. We conclude that transparency improves liquidity in the low cost, high adverse selection environment: it raises volumes and reduces bid-offer by unlocking (cheap) two-sided principal trading in the high value security. These results comport with the generally positive effect on liquidity that accompanied the introduction of TRACE in the US in the early 2000s. Finally, it is worth noting that transparency also reduces the use of agency trading in this setting. All trading was done on a principal basis in *Lemma 2*, whereas it is possible that some agency trading take place absent transparency (if the dealer utilizes either the agency-only or menu options for the high value security).

Proposition 1: With a low cost and high adverse selection, transparency increases liquidity and reduces the use of agency trading.

High cost, low adverse selection

When adverse selection is low a pooling equilibrium is possible (albeit not guaranteed). An important consequence of a pooling equilibrium is that the dealer earns differential profits on the two types of securities. The low value security is sold to the buyer at a premium to its intrinsic value, and the high value security is sold at a discount to its intrinsic value. This implies that the threshold cost above which the dealer prefers the menu over principal-only trading is now security specific. For the low value security, principal-only trading is preferred when:

$$\begin{aligned}
& \Delta_d + p * [\Delta_d + (E^*(v) - v_l)] - c * (1 - p) \\
& \geq 2(1 - q) * p * \Delta_d + q * [K + p * (\Delta_d + (E^*(v) - v_l)) - c * (1 - p)] \rightarrow \\
& c \leq c' + p * (E^*(v) - v_l) / (1 - p) \tag{22}
\end{aligned}$$

The constraint in [22] is looser than that computed with transparency (in [13] above); principal-only trading is optimal over a wider range of costs.

In contrast, principal trading is less attractive for the high value security. The dealer prefers principal-only trading when:

$$c \leq c' - p * (v_h - E^*(v)) / (1 - p) \tag{23}$$

The gap between the constraints in [22] and [23] raises the intriguing possibility that the dealer provides differential liquidity in the two securities. It operates as if its inventory cost is “low” when presented the low value security but as if its inventory cost is “high” when presented the high value security. Of course, this would change the relative proportions of securities in dealer inventory; the low value security would be overrepresented because the dealer only executes a principal trade in the high value security when the seller experienced the large liquidity shock. Specifically, the proportion of low value securities is equal to θ_l :

$$\theta_l = \frac{\theta}{[\theta + q * (1 - \theta)]} > \theta \tag{24}$$

To be an equilibrium, it must be the case that the dealer is still willing to sell both securities at the (lower) conditional expected value that reflects the increased proportion

of low value securities. In other words, adverse selection is worse for an inventory cost in this range, and pooling in this range of costs is only viable if:

$$c \geq \max[\theta_l * (v_h - v_l) - \Delta_d, \theta_l * (v_h - v_l) - c'] = c''(\theta_l) \quad [25]$$

We now further restrict the range of inventory costs we consider. First, we determine a lower bound by combining the constraints from [23] and [25]:

$$c \geq \max[c''(\theta_l), c' - p * (v_h - E^*(v))/(1 - p)] = c_- \quad [26]$$

Second, we determine an upper bound based on [22] and [15] (that the cost is not so high as to deter all principal trading):

$$c \leq \min[c' + p * \frac{(E^*(v) - v_l)}{1 - p}, \Delta_u] = c^- \quad [28]$$

For $c \in [c_-, c^-]$, the dealer buys all low value securities on a principal basis but offers the menu of agency and principal trading when presented the high value security. The buyer purchases all securities in inventory at a price of $\theta_l * v_l + (1 - \theta_l) * v_h$, and the high value security at $v_h + \Delta_d$ when the seller chose agency trading. This is a “hybrid” equilibrium, with “low” liquidity (i.e., a menu of agency and expensive principal trading) for the high value security and “high” liquidity (all trades executed using cheap principal trading) for the low value security.

We compare trading volume and bid-offer in this equilibrium to the transparent equilibrium in *Lemma 3* (see the Appendix for a full characterization of this pooling equilibrium). In the pooling equilibrium, volume reflects the fact that the dealer buys all low value securities and high value securities that are paired with a large liquidity shock

on a principal basis, and sells these positions when a buyer arrives. Further, the dealer (sometimes) matches buyers and sellers of the high value security when the seller experiences the small liquidity shock:

$$\text{Pooling volume} = \theta(1 + p) + (1 - \theta) * (q - pq + 2p) \quad [29]$$

We compare this to volume in *Lemma 3*. Total volumes are higher without transparency when:

$$\begin{aligned} \theta(1 + p) + (1 - \theta) * (q - pq + 2p) > q + p(2 - q) &\rightarrow \\ 1 > q + (1 - q) * p & \quad [30] \end{aligned}$$

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

The bid-offer on principal trades in this pooling equilibrium varies by security type. A low value security starts with a low bid offer ($2 * \Delta_d$) that is adjusted upwards due to adverse selection. A high value security starts with a high bid offer (reflecting the large liquidity shock) that is adjusted downwards, both to induce separation on the part of the seller, and because of adverse selection. The average bid-offer reflects the weight of

each security type in the dealer's inventory, where the two effects related to adverse selection offset:

$$\text{Pooling principal bid - offer} = \Delta_d + \theta_l * \Delta_d + (1 - \theta_l) * K \quad [31]$$

This implies that transparency always increases the average bid-offer on principal trades:

$$\Delta_d + \theta_l * \Delta_d + (1 - \theta_l) * K < K + \Delta_d \quad \rightarrow$$

$$\Delta_d < K \quad [32]$$

This is true for any $p < 1$. However, the magnitude of the increase depends on the probability of an agency match because K declines as the probability of a match increases. The immediacy offered by a principal trade must come at a better price to induce separation when a match is likely.

Both the changes in volumes and the increase in bid-offer for some types of transactions are driven by a shift in trading protocols: transparency increases the proportion of trades done in the agency trading protocol. Importantly, the magnitude of the effect is not uniform across transactions. Trades that are more difficult to match via the agency protocol that suffer the most from transparency.

Proposition 2: When inventory cost is high and adverse selection is low, transparency can reduce volume, increase the average bid-offer of principal trades, and raise the proportion of agency trades, particularly for trades that are difficult to match.

2.4 Interpretation and Predictions

By comparing *Proposition 1* and *Proposition 2*, we see that the effect of transparency reverses when the cost of inventory rises and the degree of adverse selection declines. With low cost and high adverse selection, which we believe describes the pre-GFC market, transparency leads to greater volume, lower bid-offer on principal trades, and more principal trading. With high cost and low adverse selection, which we believe describes the current market, transparency leads to lower volume, higher bid-offer on principal trades, and greater use of agency trading. A small amount of opacity / adverse selection is good for liquidity: it allows dealers to circumvent their elevated cost of inventory, which is welfare-improving for both dealers and investors. This leads to our testable predictions:

H1: In the current market environment, transparency increases bid-offer spreads for trades are difficult to match.

H2: In the current market environment, transparency increases the proportion of agency trades.

3. Constructing a database of European corporate bond transactions

3.1. Transactions reporting in Europe

To test these predictions, we need a data set of corporate bond transactions with differential transparency in a modern setting. Transparency in the US market has not changed since TRACE was introduced over two decades ago, and so we turn to Europe, where the Markets in Financial Instruments Directive (MiFiDII) rules require that

corporate bonds transactions are reported with a unique bond identifier, an exact execution timestamp, price and quantity.¹¹

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 many different reporting venues. This complicates the process of collecting, cleaning and aggregating transactions.

Transactions executed on online platforms (e.g., Tradeweb or MarketAxess) are reported by the respective platform, while OTC voice transactions are disclosed through an Approved Publication Arrangement (APA), which acts as the reporting entity on behalf of market-makers.¹² Each market-maker has one unique APA, which publishes all its voice transactions. Most leading electronic platforms also operate a separate and independent APA – e.g., Tradeweb and Tradeweb APA; Bloomberg and Bloomberg APA. For simplicity, we refer to both electronic platforms and APAs as “venues”.

We first collect data from 14 trading venues (for a detailed list refer to *Table A 1* in the Appendix). We then aggregate and clean the data (e.g., remove duplicates, reversals and amendments, etc.). We focus on euro-denominated investment grade (IG) corporate bonds executed between November 2022 – September 2023.¹³ For each transaction, we obtain the exact execution and reporting timestamp (which as we will see below are often

¹¹ For more details, refer to this [report](#) by ESMA.

¹² Data are made available on public websites, for examples see <https://www.apa.tradeweb.com/>; <https://www.bloomberg.com/professional/product/apae/>

¹³ To create this dataset we collaborated with *Propellant.digital*.

different), the cash price, the size, the venue on which the transaction was executed, the Market Identifier Code (MIC) of the venue. Importantly for our analysis, our database allows us to identify in which jurisdiction (EU or UK) each trade was executed. *Table 1* shows a toy example of the dataset we constructed.

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

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

We were able to match between 85% and 90% (by count and by volume) of the Barclays RFQs to the transaction dataset. In conversations with the trading desk we have verified that the majority of the unmatched RFQs were executed on dealer-to-dealer

electronic venues, which are not part of transaction dataset.¹⁴ Further, while we don't have a precise estimate of the size of the wholesale corporate bond market in Europe, TRACE estimates¹⁵ show that during the same time period, dealer-to-dealer activity in the US constituted c.15% of total volumes, which aligns with our matching rate. These tests give us confidence that the dataset we have constructed captures close to 100% of the dealer-to-client European corporate bond trades over the relevant period.

3.2. Measuring transaction costs – imputed round-trip cost (IRC)

We measure transaction costs using the imputed round-trip cost (IRC) (Feldhütter, (2012); Kargar, et al., (2021)).¹⁶ To construct the IRC, we first identify pairs of round-trip trades. A round-trip trade consists of two matched trades in the same bond with the same trade size that are executed as close as possible to each other but at different prices.¹⁷ When a trade has more than one match, we use the match closest in time. On an intuitive level, the goal of our methodology is to impute the direction of trades and, in so doing, identify a sale from an investor to a market-maker, and the subsequent purchase by another investor from the same market-maker, or vice versa. Then, for each round-trip

¹⁴ 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.

¹⁵ TRACE explicitly differentiates between dealer-to-client and dealer-to-dealer volumes. Nothing in the existing literature suggests the share of dealer-to-dealer volumes vary between the US and Europe.

¹⁶ 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 (e.g., Amihud's (2002) price impact or Roll's (1984) autocovariance in price returns) produce noisy estimates when applied at the transaction level.

¹⁷ Goldstein & Hotchkiss (2020) refer to these types of round-trips as “paired round-trips”. In other methodologies, trades are matched with more than trade in the same bond up to the face amount of the initial buy trade.

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. We remove zero-cost round-trips and *IRC* values above the 95th percentile of the distribution to ensure that our results are not polluted by extreme values.

Following the recent literature (Kargar, et al., (2021)), we explicitly differentiate between agency and principal round-trips. We identify agency round-trips as trades executed within 15 minutes, whereas principal round-trips are trades executed more than 15 minutes apart.¹⁸

Our final sample contains c.666K observations (round-trips), of which c.630K are principal round-trips and c.36K are agency round-trips. Further, c.425K roundtrips were executed in the EU, while the remaining c.221K trades took place in the UK.

3.3. Summary statistics of imputed round-trip costs

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,

¹⁸ In robustness checks, we used different thresholds (5minutes or 10 minutes) and obtained similar results.

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.

Robustness

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

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

Second, we aggregate the *IRC* to the bond-month level and compare the estimates to LCS. LCS is a commercially available measure of transaction cost computed using quotes from the Barclays trading desk. It follows the methodology by Konstantinovskiy, Yuen Ng, and Phelps (2016). LCS measures the transaction cost for an institutional-size

trade, expressed as a percentage of the bond's price (hence higher LCS signifies lower liquidity). *IRC* closely tracks LCS (*Figure 2*).

4. Exogenous variation in transactions reporting

The trade reporting rules in MiFiDII are more complex than those in the US. They 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. In practice, the reporting of virtually all transactions that qualify for a delay are in fact delayed for the full four weeks.

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 3*).¹⁹ ESMA makes a liquidity

¹⁹ 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](#).

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.²⁰ Each year, ESMA also sets two global trade size thresholds.²¹ 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 override the bond liquidity classification recommended by ESMA or to extend further the reporting deferral.

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*”

²⁰ 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.

²¹ 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.

on the execution of all other components”.²² 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.²³

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), even though 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.

4.2. Exogenous variation N.1 - Brexit

Before Brexit, ESMA had sole responsibility for liquidity assessments and making recommendations for transaction deferrals. Each bond-quarter had a unique liquidity classification (liquid or illiquid) and unique thresholds separating small from large

²² Refer to ESMA’s [guidelines](#) on the treatment of TPACs.

²³ 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.

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.²⁴

To demonstrate, in *Figure 4* we plot the percentage of transactions reported in real-time for bond-quarters classified as liquid by ESMA. We bucket transactions based on trade size and show the respective number for each trade bucket. We would expect transactions in liquid bonds below the size thresholds to be reported real-time since the reporting cannot be deferred. However, within each size bucket, we find that a substantial

²⁴ 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.

percentage of these transactions are reported with a delay. For instance, 28% of the transactions smaller than €500K are reported with a delay (dark blue bars in *Figure 4*). 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 most 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

Brexit also restricted which legal entities investors trade with. Before Brexit, most leading trading venues (e.g., Tradeweb, MarketAxess) served all 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 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...”.²⁵ Similarly, in order to be eligible to trade with the EU entity of Tradeweb, an investor must be “...*authorised under MiFID II, a credit institution authorised under EU Directive 2013/36/EU or an EU branch of a non-EU investment firm or credit institution...*”.²⁶ 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 an EU entity and vice versa. In other words, the variation in reporting delays driven by the jurisdiction effect is exogenous; reporting is imposed on investors based on their location.

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

On the 19th October 2022, ESMA announced that it would not publish the next-quarter bond liquidity assessment due to a data quality issue.²⁷ In accordance with the MiFiDII playbook, all bonds for which no liquidity assessment had been published were deemed illiquid²⁸ from 16th November 2022 until the application of the next liquidity assessment on the 16th February 2023. Therefore, all transactions in these bonds qualified for a

²⁵ Tradeweb UK’s [Rulebook](#).

²⁶ Tradeweb EU’s [Rulebook](#).

²⁷ The press release can be found [here](#).

²⁸ This is in line with Q&A 10 of section 4 of the MiFID II transparency Q&As.

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 outlining how this affected the reporting requirement.

In the EU, the number of transactions reported with transparency decreased sharply on the 16th November 2022 and subsequently recovered on the 16th February 2023, when the next regular publication period began and the new reporting rules applied (Panel A, *Figure 5*). 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 5*). Again, there was no corresponding effect in the UK. Our analysis shows that the UK venues most likely applied the last published classification (i.e., the classification used for the period 16th August 2022 to 16th November 2022) for bonds issued before November 16th and reported all bonds issued during the “no publication” period with transparency.

This “no publication” event generated two additional sources of variation:

1. **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 can study the effect of transparency on bond liquidity twice, as treated bonds both enter and exit the “no publication” period.

2. **Across jurisdictions:** the reporting of some bonds changed in the EU, but it did not change in the UK. This applied to bonds aged between 6 months and 2 years – i.e., the difference between the light blue and dark blue lines in Panel B of *Figure 5*.

4.4. Empirical Design

We use this exogenous variation in reporting in two ways. In *Section 5*, we exploit 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. This allows us to compare transaction costs for “treatment” and “control” bonds, before and after transparency changed, using a difference-in-differences framework. In *Section 6*, we study how the effect of transparency varies depending on the probability of finding a match. To do this robustly, we need a larger sample size, which is why we pool all the round-trip trades, and exploit the total variation in reporting at the transaction-level 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.

5. Difference-in-differences (DID) estimates

5.1. 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 can 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.²⁹ 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 A 2* we verify that transactions in treated and control bonds were in fact reported as expected. We calculate that 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

²⁹ 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.

transparency before the 16th February 2023 and 98.4% were reported with transparency afterwards. Close to 90% of transactions in control bonds were reported with transparency in all cases.

The DID approach takes a treated bond before and after the treatment and compares its outcome 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 study the effect of transparency on two different outcomes: the share of agency trading and transaction costs.

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 all specifications (see below *Table 4*), conditional on observable characteristics, the difference between treated and controls is economically small and/or statistically insignificant, which we would expect given that only a slight difference in age separates the two categories of bonds. Finally, in the Appendix we verify that the parallel trends assumption holds both for the share of agency trades (*Figure A 3*) and for transaction costs (*Figure A 4*).

5.2. Transparency increases the share of agency trading

We compute the proportion of agency trades for treated and control bonds shortly before and after they enter and exit the “no publication” period (*Figure 6*). Control bonds do not change reporting status, as they are always reported with transparency, and hence, the proportion of agency trading for these bonds remains unchanged both as they enter

and exit the quasi-natural experiment window. However, as treated bonds change their reporting from transparency to no transparency (i.e., as they enter the “no publication” event) the proportion of agency trading drops by roughly half (from 10% to 5.9%). We obtain the mirror image on the other side of the event window, as treated bonds exit. As treated bonds change their reporting from no transparency to transparency, their proportion of agency trading increases from 8% to 15.4%. These results align with the predictions of our theoretical model.

5.3. Transparency increases transaction costs

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

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

(Model 1)

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 estimate two sets of conceptually equivalent DID regressions – (1) as bonds enter and (2) as bonds exit the “no publication” window (**Table 4**).

To isolate the effect of transparency on transaction costs from the effect of other bond characteristics which independently drive transaction costs, we control for variables at the round-trip level (collected in the vector $X_{i,j,t}$), as well as for time-varying bond-level characteristics (collected in the vector $Z_{j,t}$). Round-trip level controls include the number

of days it takes to close a position³⁰, an electronic trade dummy³¹ and a package trade dummy. Time-varying bond characteristics include bond age (years since issuance), the logarithm of amount outstanding, remaining years to maturity and a credit rating dummy. Finally, we also include bond ($Treated_j$) and date ($Post_t$) fixed effects to account for any (potentially unobservable) factors that could affect our results.

The transaction cost of principal trades in treated bonds declines by 6bp when they enter the no-publication period on the 16th November 2022 ($\beta_1 < 0$). Similarly, as treated bonds exit the no-publication period on the 16th February 2023, the bid-offer on principal trades increases by 3 bp ($\beta_1 > 0$). In percentage terms, this corresponds to an increase in bid-offer by 6%-10%. Both results are statistically significant and support the conclusion of our empirical model that transparency can increase the bid-offer of principal trades.

6. Pooled regression estimates

6.1. Econometric model

As a next step, we pool all sources of exogenous variation and estimate regression models at the transaction-level. Identification comes from variation in the transaction costs of bonds which have a different reporting schedule in different jurisdictions. This increases the sample size, allowing us to examine how the effect of transparency varies across trades depending on how difficult they are to match.

³⁰ *Figure 1* 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.

³¹ Anecdotal evidence suggests that electronic venues have better reporting discipline and commit fewer reporting errors than APAs, which report voice transactions.

We proxy difficult-to-match trades in two ways: by size and by age. Corporate bonds trade infrequently, and typically have low turnover. One feature of the corporate bond market is that the number of trades and the value of trades are not uniformly distributed by trade size. For example, trade sizes smaller than €500K account for 80% of the observations but only 20% of the total notional traded. Conversely, large trades account for a small number of the total number of observations but generate most of the volume (*Figure A 3*). This is why it is substantially more difficult for a market-maker to find a match for a €2M position, compared to a €500K position. Also, bonds trade very frequently shortly after they are issued, after which their liquidity sharply declines – for example average monthly turnover decreases from 25% for newly issued bonds to less than 5% for bonds issued more than five years ago (*Figure A 4*). Hence, finding a match for a large position and/or a position in an aged bond is more difficult.

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

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

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). Hence β_1 gives the difference between the transaction-cost of roundtrips reported with and without transparency. $Difficult\ to\ Match_{i,j,k,t}$ is either trade size buckets or bond age. In the former case, we define four buckets: $\leq €500K$, $(€500K-€1M]$, $(€1M-€2M]$ and $(€2M-M3.5M]$, where the

\leq €500K category is our reference bucket – hence the effect of transparency for trades smaller than €500K is given by β_1 and the effect for trades larger than €500K is given by the sum of β_1 and β_2 . In the latter, we use a continuous measure of bond age.

We include the same set of round-trip level controls ($X_{i,j,k,t}$) and time-varying bond controls ($Z_{j,t}$) as in our baseline DID specification. We also add bond (λ_j), jurisdiction (δ_k) and date (γ_t) fixed effects (for the first leg of the round-trip) to account for any (potentially unobservable) factors that could affect our results. We estimate the model for principal trades, and use agency trades in robustness checks.

6.2. Transparency increases transaction costs for trades that are difficult to match

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*). In *Figure 7* 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.20% increase. Weighing the effect of transparency for each trade size bucket shown on *Figure 7* by its contribution to total volumes, we calculate that on average, the effect of transparency is a 6% increase in transaction cost for principal trades, which aligns with our baseline DID results (*Table 4*).

We obtain similar conclusions when we look at the interaction between transparency and bond age (column (2), *Table 5*). To demonstrate that these effects apply

independently of those of trade size, we also retain the interaction terms with trade size buckets in the regression specification. As a bond ages, transparency has a more deleterious effect on liquidity ($\beta_2 > 0$).

6.3. Robustness

A large portion of our round-trips are not completed on the same day, which is why it is possible that bid-offer spreads capture changes in market conditions which occur while the market maker is looking for the other side of the trade. For example, the market maker buys a large trade in a bond reported with transparency on a day when volatility is low, but offloads the position on a different day when volatility is high. Differences in bid-offer compared to similar trades reported without transparency might be related to changes in volatility, and not to transparency. In column (1) of *Table 6* we include two-way date-fixed effects (i.e., a date dummy for the first and second leg of the round-trip) to account for the possibility that the first and the second leg of the round-trip might be executed under different market conditions, and verify that our results remain unchanged.

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 column (2) of *Table 6* we include venue fixed effects and find very similar results to the results reported 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 (column (3)). We obtain similar results with specifications that limit our sample to round-trips where both legs are in the same jurisdiction ((column (4), *Table 6*).

As a final robustness check, we estimated the same model for agency trades and find that transparency has no impact on bid-offer spreads, which is aligned with the intuition of our theoretical model (columns (5) and (6), *Table 6*). Agency trades are pre-negotiated, meaning that at the time the first leg of the round-trip was reported, the market maker had found a matching buyer. Hence, whether agency round-trips are initiated with transparency or not should have no implications for transaction costs.

7. Discussion and policy implications

Revisions to transparency have been proposed in a number of jurisdictions and across asset classes. In drafting these proposals, regulators have cited the existing literature based on TRACE data, which concludes that transparency unequivocally improves liquidity.³² An important contribution of our work is to challenge this assertion. We conclude that the effect of transparency has changed in response to the increase in dealer inventory costs and the decline in the potential for adverse selection. It is now heterogeneous, dependent on the combination of trade and bond characteristics. Liquidity in transactions that are more difficult to match in the agency trading protocol suffers from

³² Edwards, Harris, & Piwowar (2007); Bessembinder, Maxwell, & Venkataraman (2006); Goldstein, Hotchkiss, & Sirri (2006)).

transparency, which both pushes more volume into (uncertain) agency trading and increases the bid-offer of principal trades.

We consider the implications of our analysis for the extant proposals on transparency. Starting with Europe, regulators in both the EU and in the UK have recently published proposals to amend the existing framework for reporting corporate bond transactions, in an effort to enhance market data transparency and reduce fragmentation.³³ Although the specific provisions and technical details differ slightly, the overarching goal in both jurisdictions is to implement faster disclosures for corporate bond trades and establish a consolidated tape as a single reference source of information for prices and volume of traded bonds. We estimate that the new rules will increase the fraction of European corporate bond transactions reported in real-time from 8% to c.80%.³⁴

Our work supports maintaining some version of the current sliding scale for reporting. Large trades make up a relatively small fraction of total trades but a significant majority of total volume. Further, these trades are likely to be more important in times of market or economic turmoil, when investors may have more significant security-specific liquidity needs. That said, the optimal size thresholds and reporting delays are unclear; we can only study the status quo. One recommendation we make is to implement changes

³³ 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.

³⁴ 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](#).

slowly, in stages, allowing for study of the relevant tradeoffs. For example, regulators may be willing to accept some decline in average liquidity in exchange for improved transparency, which may have benefits away from liquidity (e.g., for the primary market). In addition, investors and dealers may adapt to a new transparency regime in ways that alter its effect on transparency; our quasi-natural experiment was short-lived, and thus was unlikely to evince material changes in strategy or use of new protocols.

Regulators in the US are also considering new transparency rules. First, FINRA is considering shrinking the 15-minute reporting window for corporate bond trades. Our results suggest this may be counterproductive, at least with regards to large trades. Smaller trades are increasingly executed in electronic or “PT” (portfolio trade) form; in practice these come with real-time reporting, and thus a shorter window is unlikely to have a meaningful effect on their liquidity characteristics.

Similarly, there is discussion of extending the public reporting of on-the-run US Treasury trades to transactions in off-the-run securities. Our results regulators should proceed with caution on this proposal. There is an extensive literature documenting how the post-crisis regulatory reforms, particularly the supplemental leverage ratio, raise the cost of intermediation of Treasuries, leading to a notable, and concerning, decline in Treasury liquidity. This is most apparent in off-the-run securities, in keeping with our framework, which suggests that securities which are more difficult to “match” suffer the most from an increase in inventory cost. At the same time, adverse selection in this market is quite low: there is no dearth of information about the shape of the yield curve, and there are no security- or issuer-specific nuances to consider when valuing US

Treasuries. Our framework implies that transparency would only worsen liquidity in these transactions.

Our results also suggest that any transparency regime has significant welfare implications. Greater transparency likely hurts investors who have a greater need to trade individual securities in size, such as institutional investors which implement traditional fundamental-based strategies. They are likely to face significant liquidity needs in specific securities as the market (and their views) change. We expect it has a limited effect on investors who trade individual securities in small size, such as retail investors and the new breed of “quant credit” investors. This latter category has expanded recently, as investors take advantage of new trading protocols, notably portfolio trading, to implement factor-based investment strategies over a large number of securities. These welfare implications are interesting avenues for further study.

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

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.
- An, Y., & Zheng, Z. (2022). *Immediacy Provision and Matchmaking*. Retrieved from Available at SSRN: <https://ssrn.com/abstract=2868280>
- Asquith, P., Covert, R., & Pathak, P. (2019). The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market. *NBER Working Paper 19417*.
- Back, K., Liu, R., & Tegui, 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 Transaction Costs of Index CDSs. *The Journal of Finance*, 75(5), 2719-2763.
- Dick-Nielsen, J., & Rossi, M. (2018). The Cost of Immediacy for Corporate Bonds. *The Review of Financial Studies*, 32(1), 1-41.
- Duffie, D., Dworczak, P., & Zhu, H. (2017). Benchmarks in Search Markets. *The Journal of Finance*, 72(5), 1983-2044.
- Duffie, D., Gârleanu, N., & Pedersen, L. (2005). Over-the-Counter Markets. *Econometrica*, 73(6), 1815-1847.
- Edwards, K., Harris, E., & Piwowar, M. (2007). Corporate Bond Market Transaction Costs and Transparency. *The Journal of Finance*, 62(3), 1421-1451.

- Feldhütter, P. (2012). The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures. *The Review of Financial Studies*, 25(4), 1155-1206.
- Glebkin, S., Yueshen, B., & Shen, J. (2022). Simultaneous Multilateral Search. *The Review of Financial Studies*, 571-614.
- Goldberg, J., & Nozawa, Y. (2020). Liquidity Supply in the Corporate Bond Market. *Journal of Finance*.
- Goldstein, M., & Hotchkiss, E. (2020). Providing liquidity in an illiquid market: Dealer behavior in US corporate bonds. *Journal of Financial Economics*, 135(1), 16-40.
- Goldstein, M., Hotchkiss, E., & Sirri, E. (2006). Transparency and Liquidity: A Controlled Experiment on Corporate Bonds. *The Review of Financial Studies*, 20(2), 235-273.
- Hagströmer, B. (2021). Bias in the effective bid-ask spread. *Journal of Financial Economics*, 142(1), 314-337.
- Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.-O., & Zúñiga, D. (2021). Corporate Bond Liquidity during the COVID-19 Crisis. *The Review of Financial Studies*, 34(11), 5352-5401.
- Konstantinovskiy, 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.
- Saar, G., Sun, J., Yang, R., & Zhu, H. (2022). *From Market Making to Matchmaking: Does Bank Regulation Harm Market Liquidity?* Retrieved from Available at SSRN: <https://ssrn.com/abstract=3399063> or <http://dx.doi.org/10.2139/ssrn.3399063>
- Todorova, Z., & Diaz, A. (2023). There's still a role for voice. *Barclays research series*.
- Vairo, M., & Dworzak, 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: IRC: agency vs. principal trades

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




Figure 2: IRC vs. LCS

The figure compares the monthly average weighted IRC and LCS. LCS is an alternative measure of transaction-costs computed using quotes from the Barclays credit trading desk.




Figure 3: Post-trade reporting rules

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

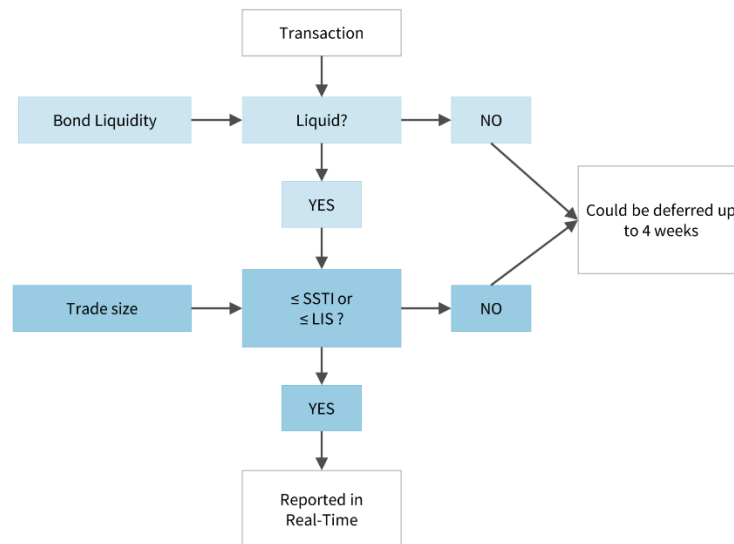


Figure 4: Variation in transactions reporting N.1 – Brexit effect

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.

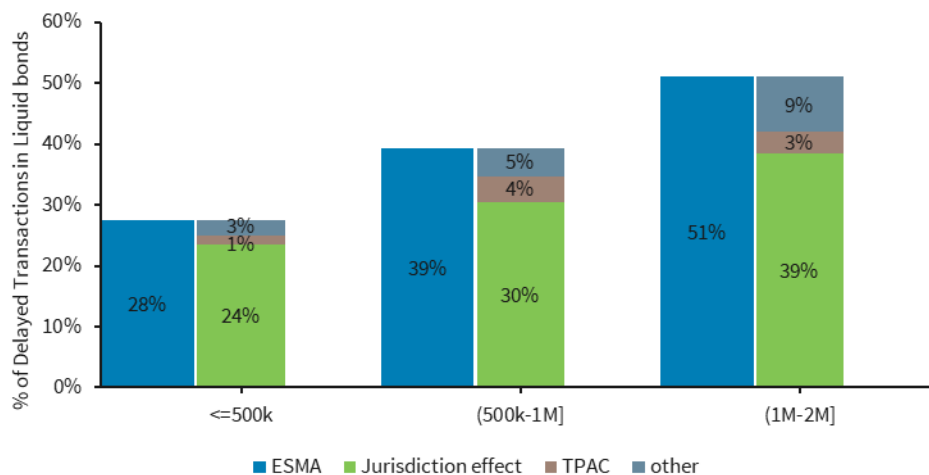



Figure 5: Variation in transactions reporting N.2 – shutting down transparency in the EU

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 6: Difference-in-differences – agency trading

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





Figure 7: The effect of transparency by trade size

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



List of Tables

Table 1: A snapshot of the data

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

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

Table 2: Bond characteristics – EU vs. UK

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

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

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

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

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

Table 4: Difference-in-differences regressions – the effect of transparency on transaction costs

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

The table reports difference-in-differences regressions of Imputed Round-trip Cost ($IRC_{i,j,t}$) executed in the EU on a Treated dummy ($Treated_j$), Post dummy ($Post_t$) and their interaction term ($Treated_{i,j} \times Post_t$) and a set of controls. Regressions include the following controls at the round-trip level ($X_{i,j,t}$): number of days to close a position, electronic trade dummy and a package transaction dummy. Time-varying controls include ($Z_{j,t}$): 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 column (1) based on data from the 1st November 2022 to the 30th November 2022; results in column (2) are based on data from the 15th Jan 2023 to the 15th March 2023. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

	IRC, bp	
	(1) Entering the “no publication” period	(2) Exiting the “no publication” period
Treated × Post	-6.19*** (-2.50)	3.23*** (2.85)
Treated	1.86*** (3.00)	0.11 (0.14)
Post	1.10 (0.86)	1.07 (1.49)
Round-trip level controls	YES	YES
Bond-date level controls	YES	YES
Size FE	YES	YES
Round-Trips Observations	1 Nov 2022- 30 Nov 2022 3,157 trades	15 Jan 2023-15 March 2023 11,875 trades

Table 5: Pooled regressions – the effect of transparency on transaction costs varies depending on how difficult it is to match trades

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

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

Table 6: Pooled regressions – 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 ($Transparency_{i,j,k,t} \times Size\ Bucket_s$) 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 ($Z_{j,t}$): 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 (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] (trades in the <€500K are the reference category). Column (1) adds date fixed effects for both legs of the round-trip; column (2) adds trading venue fixed effects; column (3) limits the sample to round-trips where both legs are in the same jurisdiction; column (4) includes jurisdiction-date fixed effects; columns (5) and (6) use agency round-trips. T-stats in parentheses. Significance at the 1 %, 5 % and 10 % statistical level is denoted by ***, **, and * respectively.

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

Model Appendix

Lemma 1 establishes the necessary equilibrium decision rules for the seller and buyer in a trembling hand perfect equilibrium. This reduces computing the equilibria to finding the profit maximizing strategy for the dealer, conditional on these decision rules.

In each of the equilibria that we consider, we define a range of inventory cost such that determines the dealer's optimal strategy. For example, in *Lemma 1*, the "low" cost regime implies that the dealer maximizes its profits with a principal-only bid. Therefore, the uniqueness of the resulting equilibrium follows. The same holds for the equilibrium in *Lemma 2*. We need only verify the equilibrium trading volume in that case:

$$Volumes = q(1 + p) + (1 - q) * 2 * p = 2p + q - pq \quad [A1]$$

Low cost, high adverse selection Equilibrium and Proof of Proposition 1

As detailed in the text, there are two possible decision rules for the buyer when adverse selection is high. It may only pay a price associated with the low value security, or it may be willing to pay a higher price if it experiences the large liquidity shock. The latter is only possible if the dealer only sells the high value security at that price, which requires that the dealer prefer the certainty of $v_l + \Delta_d$ when it owns the low value security over $v_h + \Delta_u$, but only trading with probability q . This is true when:

$$v_l + \Delta_d \geq q * (v_h + \Delta_u) + (1 - q) * (v_l - c) \rightarrow$$

$$c \geq \frac{q}{1-q} * (v_h - v_l) - c' \quad [A2]$$

If the inventory cost is high relative to the probability of the large liquidity shock, the dealer does not “cheat” and try to sell the low value security at the high price. Note that [A2] depends on the distribution of liquidity shocks, whereas the degree of adverse selection depends on the distribution of security types, so it is possible that [A2] is met even when adverse selection is high (and cost is low). If [A2] is met, then the buyer must be willing to trade at the higher price when it experiences the large liquidity shock, as that maximizes the probability that the buyer transacts.

Under either decision rule, trading in the low value security follows *Lemma 2*. The dealer uses agency-only, principal-only, or a menu when quoting the high value security; the choice depends on the inventory cost, and the thresholds depend on which decision rule the buyer uses. Assume first the buyer trades when $A \leq v_l + \Delta$. We compare the profits from each trading protocol. The agency only option generates a profit of $2p\Delta_d$. The optimal principal-only bid is $v_h - \Delta_d$; it is preferred to $v_h - \Delta_u$ because the inventory cost is “low”:

$$\begin{aligned} \Delta_d - c &\geq q(\Delta_u - c) \rightarrow \\ c &\leq c' \end{aligned} \quad [A3]$$

Therefore, the principal-only option generates profit of $\Delta_d - c$. The menu generates profit of $2(1 - q)p\Delta_d + q(K - c)$. The agency-only option is more profitable than the principal-only option when $c > \Delta_d(1 - 2p)$. It is more profitable than the menu when $c > (1 - p)\Delta_u - p\Delta_d$. We use a strict inequality because the agency-only option has a

lower volume than either of the other options. Therefore, the dealer chooses the agency-only option when:

$$c > \max [\Delta_d(1 - 2p), (1 - p)\Delta_u - p\Delta_d] \quad [A4]$$

The menu is chosen over the principal-only option when:

$$\begin{aligned} \Delta_d - c < 2(1 - q)p\Delta_d + q(K - c) &\rightarrow \\ c > (1 - p)c' - p\Delta_d & \quad [A5] \end{aligned}$$

Therefore, the dealer chooses the menu when:

$$c \in ((1 - p)c' - p\Delta_d, \max [\Delta_d(1 - 2p), (1 - p)\Delta_u - p\Delta_d]] \quad [A6]$$

Otherwise, it chooses the principal-only option.

If instead [A2] is satisfied and the buyer will purchase the high value security, the profits associated with principal trades in that security are higher, increasing the cost ranges over which the dealer uses principal-only and the menu trading. Principal-only trading in the high value security generates profits of $\Delta_d + pq * \Delta_u - (1 - pq) * c$. The menu generates profits of $2(1 - q)p\Delta_d + q(K + pq * \Delta_u - (1 - pq) * c)$

Agency-only is now preferred when:

$$c > \max\{[\Delta_d(1 - 2p) + pq\Delta_u]/(1 - pq), [(1 - p)\Delta_u - p(\Delta_d - q\Delta_u)]/(1 - pq)\} \quad [A7]$$

The menu is the preferred option when:

$$\begin{aligned} c \in \{(1 - 2p + pq)c'/(1 - pq), \\ \max\{[\Delta_d(1 - 2p) + pq\Delta_u]/(1 - pq), (\Delta_u - p\Delta_d)/(1 - pq)\} \} \quad [A8] \end{aligned}$$

Otherwise, the dealer prefers principal-only trading. These ranges are uninteresting except to note that agency-only trading is preferred over a narrower range, because the profitability of both principal-only trading and the menu are higher if there is some two-sided trading in the high value security. We can now characterize the equilibrium.

The unique low cost, high adverse selection, non-transparent equilibrium is:

- a) *The dealer buys all low value securities from the seller on a principal basis at $v_l - \Delta_d$, and sells them to the buyer (when it arrives), at $v_l + \Delta_d$;*
- b) *For the high value security, the dealer engages in either agency-only trading, one-sided principal-only trading, where it holds the security in inventory, or a menu of the two, using K as in [11], depending on the inventory cost, using the relevant ranges (which depend on whether [A2] is satisfied). In the latter case, the seller chooses immediacy when it faces a large liquidity shock and the agency protocol when it faces a low liquidity shock;*
- c) *If [A2] is satisfied the dealer sells the high value security to the buyer at $v_h + \Delta_u$ when it experiences the large liquidity shock;*
- d) *Otherwise there is no two-sided principal trading in the high value security.*

The proof of *Proposition 1* follows immediately. Volume declines because there is no two-sided principal trading in the high value security when the buyer experiences the small liquidity shock (it may decline more than this, if [A2] is not satisfied). The average bid-offer on principal trades rises if [A2] is satisfied, because round-trip trades in the high value security have a bid-offer of $\Delta_u + \Delta_d$ (compared to $2\Delta_d$ for round-trip

trades in the low value security). Finally, for some cost ranges, the dealer uses some agency trading, which is not the case when the market is transparent. QED

High cost, low adverse selection equilibrium and Proof of Proposition 2

We demonstrate in the text that for $c \in [c_-, c^-]$ a pooling equilibrium obtains where the dealer provides the seller differential liquidity depending on the security type. The dealer bids $v_l - \Delta_d$ when the seller owns the low value security and uses the menu when the seller owns the high value security. This leads to the following equilibrium:

With $c \in [c_-, c^-]$ and no transparency, the unique equilibrium is:

- a) For the high value security the dealer offers the seller a choice of a certain principal at $v_h - K$ or an agency trade at $v_h - \Delta_d$ (which has success rate p), for K defined in [11];*
- b) For the high value security the seller chooses principal trading when it faces a large liquidity shock and agency trading when it faces a low liquidity shock;*
- c) The dealer buys all low value securities on a principal basis at $v_l - \Delta_d$;*
- d) The buyer (if it arrives) purchases all securities in dealer inventory at $\theta_l * v_l + (1 - \theta_l) * v_h + \Delta_d$;*
- e) If the seller chooses the agency trade, the buyer (if it arrives) purchases the security at $v_h + \Delta_d$.*
- f) Total transaction volume equals $\theta(1 + p) + (1 - \theta) * (q - pq + 2p)$]*
- g) Realized bid-offer on round trip principal trades equals $\Delta_d + \theta_l * \Delta_d + (1 - \theta_l) * K$.*

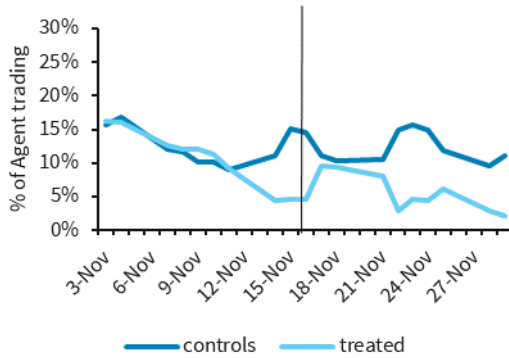
Empirical Appendix

A1. Figures

Figure A 1: Parallel trends – agency trading

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

Panel A: Entering the “no publication” period



Panel B: Exiting the “no publication” period





Figure A 2: 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




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: 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 Konstantinovskiy, V. (2023).



A2. Tables

Table A 1: List of Trading Venues

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

Table A 2: Treated vs. controls

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

	% trades reported with transparency in the EU			
	Entering the “no publication”		Exiting the “no-publication”	
	Pre (1-15 Nov 2022)	Post (16-30 Nov 2022)	Pre (15 Jan -15 Feb 2023)	Post (16 Feb -30 Mar 2023)
Controls	88.5%	76.0%	81.5%	92.3%
Treated	89.5%	0.0%	0.0%	98.4%