Incentives to Lose: Disclosure of Cover Bids in OTC Markets*

Andrey Ordin and Ruslan Sverchkov†

October 14, 2024

Abstract

We study incentives for post-trade information disclosure in the over-the-counter financial markets. In a setting where execution prices alone do not fully capture the value of the traded assets, we model decisions of investors to hide or reveal information embedded in unexecuted offers. Our model explains why investors, requesting quotes from multiple dealers in the corporate bond market, might choose to conceal the runner-up offer — the cover — from the winning dealer even though the increased informational opacity decreases dealers' incentives to win the trade and worsens their quotes. Investors conceal covers if they trade frequently, gains from trade are high, or uncertainty about bond values is low. We discuss the implications for market liquidity, fragmentation, and the design of electronic RFQ platforms.

Keywords: OTC markets, post-trade transparency, RFQ, cover bids.

*We thank Bruno Biais, Jamie Coen, Jean-Edouard Colliard, Vincent Glode, Bart Zhou Yueshen and audiences at McGill University, University of Warwick, and the Plato Partnership Conference 2024 for their helpful comments.

†Andrey Ordin (andrey.ordin@mccombs.utexas.edu) is at the McCombs School of Business, the University of Texas at Austin, and Ruslan Sverchkov (ruslan.sverchkov@wbs.ac.uk) is at Warwick Business School, the University of Warwick.

1 Introduction

The US corporate bond market, as an over-the-counter (OTC) financial market, can be characterized as informationally opaque (Duffie, 2012). Until recently, post-trade transparency was limited and information about the history of bond transactions was not easily obtained. Transparency in the market significantly improved in 2004 with the introduction of TRACE[1] which disseminates prices, quantities, and directions of all transactions after their execution to the market participants.

Yet, the OTC trade process generates more information than the execution data — this information is contained in the unexecuted offers refused during the trade. Indeed, to trade a bond, an investor usually contacts multiple dealers and asks for a price quote. This can be done either sequentially or, if the trade is conducted on an electronic platform, simultaneously via a request for quote (RFQ).[2] The investor observes all the dealers' quotes and trades with the dealer offering the best price. This price has to be disclosed to all market participants via TRACE. However, besides the winning bid, the investor also learns other bids, including the second-best bid which is colloquially called a cover. In this paper, we study the value of the additional information contained in the cover as well as the investor's incentives to disclose it after a trade.

We are motivated by the observation that investors differ in their choices to reveal or conceal the cover. On the one hand, conventionally, many investors voluntarily disclose the runner-up bid to the winning dealer. This practice is sufficiently common that it was even encoded in some electronic RFQ platforms, including MarketAxess and Tradeweb, during their design.[3] On the other hand, some investors in the market develop a reputation for

[1]The Trade Reporting and Compliance Engine (TRACE) is a platform developed by FINRA and instituted by the Securities and Exchange Commission under the Rule 6200 Series, which mandates disclosure of trade information, including price, volume and direction, about past bond transactions. The effects of TRACE are studied by Bessembinder, Maxwell and Venkataraman (2006); Edwards, Harris and Piwowar (2007); Goldstein, Hotchkiss and Sirri (2007).

[2]Electronic corporate bond trade has been steadily increasing (see, e.g., Hendershott and Madhavan, 2015; O'Hara and Zhou, 2021; Hendershott, Livdan and Schürhoff, 2021). Bloomberg Intelligence (2023) estimates that, in the beginning of 2023, the electronic activity accounted for about 38% and 33% of investment-grade and and high-yield corporate bond trading, respectively.

[3]See SIFMA (2016) that provides descriptions of protocols on various electronic bond trading platforms. See also Fermanian, Guéant and Pu (2016); Kozora, Mizrach, Peppe, Shachar and Sokobin (2020).

refusing to disclose covers ex-post, including major corporate bond traders such as PIMCO.[4] Similarly, off-platform trade leaves an opportunity for investors to conceal information from the dealers. Our goal is to characterize the conditions under which the incentives to hide information emerge and to show how these incentives affect the market.

Refusal to disclose the cover might seem puzzling at first. If additional information possessed by the investor is valuable to the dealers, they should be willing to pay more for the opportunity to gather it. For ex-ante information sharing, this observation can be formalized using the linkage principle from the auctions literature (Milgrom and Weber, 1982a). Likewise, the investor's decision to hide the cover ex-post can create an incentive against aggressive bidding: the losers and the non-participants get to learn the winning bid for free via TRACE, while the winner does not learn the exact bids of the opponents. As a result, the auction winner ends up at an informational disadvantage: he does not know the extent of the winner's curse. If this disadvantage matters, investors that conceal information should expect worse quotes from the dealers. To explain why investors might forgo the trading revenues, we propose a model in which investors trade repeatedly with dealers.

We start with a model in which a seller can affect the degree of post-trade transparency. In our setup, there is an asset with some underlying value that is common to both the seller and the buyer. The seller possesses information about this value that is not available to the buyer. For example, this value could represent the best quote that the seller can obtain from other, fringe buyers. On the other hand, the buyer has an extra private benefit from obtaining the asset in addition to the commonly shared value. This private benefit could represent the expected markup that the dealer can charge when reselling the asset to his network of clients.[5] The seller thus would like to earn information rents on the trade, while the buyer would like to realize the private benefit without paying too much for it.

Importantly, the first period of trade is followed by further interaction between the same agents, in which the seller may attempt to offer one or multiple additional assets to the buyer. The key assumption in the model is that the values of the assets across the trade

[4] See Reuters (2018) that describes how PIMCO obtained an exemption from transmission of the cover information on MarketAxess platform.

[5] Di Maggio, Kermani and Song (2017) documents that dealers charge significantly higher markups when trading with their clients as opposed to other dealers.

iterations may be related to each other. Consequently, learning something about the value of the asset traded early on is potentially useful to the buyer in the future. In the analysis, we explore the difference in what the buyer learns depending on the trading outcome in the first period and the seller's choice of post-trade transparency and, consequently, how this difference affects the incentives to win.

Each period of time, the seller has the option to trade with the fringe buyers and collect the payoff equal to their best quote. The TRACE reporting system ensures that this transaction price is disclosed to everyone. The alternative is to trade with the main buyer. In this case, if the seller has chosen not to disclose the cover, the main buyer does not get to learn how much the fringe buyers would pay. Our analysis shows that the seller has no strategic incentives to reject offers better than the fringe buyers' quote; therefore, when his offer is accepted, the main buyer learns he is paying more for the asset than the rest of the market thinks it is worth. Knowing the seller's outside option exactly would allow the buyer to maximize gains to subsequent trades, but to learn it the buyer has to lose one of the trades to the fringe. Compared to the case where covers are revealed after the trade, this tension pushes the main buyer to discount his initial bid: he trades off the value of information gained from losing against the private benefit sacrificed to learn it.

The seller decides to hide a cover if the benefit of restricting the buyer's learning is more profitable than the cost of receiving less aggressive quotes in the first trade. We show that the seller maintains opacity if the potential gains from trade are large or the uncertainty about the asset values is low. Intuitively, if the buyer's private benefit is very large, he would be extremely reluctant to risk the possibility of being rejected by the seller. Therefore, even if no information is revealed after the trade, the buyer will try to make sure the bid is good enough to be accepted, which improves the seller's expected payoff. On the other hand, when the gains from trade are small, compared to the case of disclosed covers, the buyer would significantly discount his offer. From the seller's perspective, a discounted offer is unlikely to beat the fringe buyers, so she would rather commit to disclosing the cover instead.

In our model, from the social welfare perspective, the gains from trade between the buyer and the seller are always positive, so successful trade is always desirable. However, if the seller chooses to hide the cover, the trade is more likely to break down, making it a socially sub-optimal choice. Interestingly, in both cases, with a disclosed cover and with a hidden cover, the trade breaks down only in the first period. However, it happens more often in the latter case because the buyer offers more aggressive quotes, which leads to lower welfare. In addition, because the seller prefers opacity if the gains from trade are high, the welfare is not monotone with respect to the gains from trade.

Our results highlight that liquidity in opaque OTC markets such as the US corporate bond market can be shaped by post-trade transparency choices of not only dealers but also investors requesting quotes from dealers. We show that investors might optimally choose to conceal information contained in the cover bids. As a result, dealers quote less aggressive prices and some gains from trade are forgone.

Investors' transparency choices may also influence investors' choices of trading venues, contributing to market fragmentation. Investors who would like to credibly disclose covers ex-post might prefer to trade via an electronic RFQ platform, such as MarketAxess or Tradeweb, where disclosure is encoded and automatic, while those seeking to conceal covers may opt for offline trading via phone. Indeed, for the OTC markets, in Section 4 we show that the power to commit to information disclosure is not always available even if we allow for repeated interactions, i.e., relationships, between investors and dealers. Ex-post, all investors have incentives to misreport the cover but investors with weak relationships, who participate in trade only infrequently, cannot be easily punished via repeated interactions. Our baseline results, on the other hand, imply that infrequently trading investors are precisely the ones who would benefit from committing to share ex-post information with the dealers. Therefore, one potential benefit of the electronic trading platforms is to provide commitment power to such traders.

The important question arising from our analysis is how to address the inefficient transparency choices of investors. In the past, potentially inefficient post-trade transparency choices of dealers were addressed by regulation that requires dealers to report to TRACE. Similarly, one way to improve market liquidity and consolidate the market could be to mandate disclosure of covers. Disseminating this additional post-trade information to the market participants seems straightforward. The cover price would be received by a winning dealer who can then add it to other data reported to TRACE. According to our model, such mandatory disclosure would decrease trading profits of investors but would improve the efficiency of bond allocations.[6]

In the main model, to clearly present the intuition behind our results, we consider a setup in which there are two rounds of trade and the same asset is offered for sale in both rounds. However, as we show, the model can be extended to allow more general settings with more than two rounds of trade and when the traded assets differ between periods. Our main results continue to hold in these more general setups. In addition, in the model, there is only one buyer-dealer who competes with quotes of other dealers whose best offer is represented by the common value of the asset known to the seller. Our approach allows us to avoid problems endemic in complex models for common value auctions with asymmetrically informed bidders (see, e.g., Abraham, Athey, Babaioff and Grubb, 2020): for example, we do not need to rely on refinements to select a unique equilibrium, and our model remains tractable in a setup with repeated trade. In Section 6, we discuss the conditions under which we expect our results to be mirrored in models that explicitly characterize the strategic incentives of the fringe buyers.

Related Literature.

To the best of our knowledge, our paper is the first to study information contained in covering bids or covers — unexecuted offers refused during a trade. In the context of the corporate bond market, we explore private incentives for ex-post disclosure of covers and show how such disclosure affects market liquidity. We thus, contribute to several strands of the literature: on post-trade transparency, on trade in OTC markets, and on information design in auctions.

Prior literature on post-trade transparency considers the effects of disclosing data, either order flow or pricing, about executed offers. We complement this literature by studying the disclosure of the additional information contained in unexecuted offers. Theoretical papers on post-trade transparency include Madhavan (1995); Pagano and Röell (1996); Kakhbod and Song (2020); Vairo and Dworczak (2023). In the early models of Madhavan (1995) and Pagano and Röell (1996), uninformed market makers are trading with a mix of informed

[6]Note that the electronic platforms have to take into account the same trade-off when designing the trading mechanisms that they offer to their clients.

and noise liquidity traders.[7] Transparency is interpreted as a degree to which the order flow is visible to competing dealers. The more precise information about the order flow helps dealers detect informed trading and, thus, offer better pricing for uniformed traders. Madhavan (1995) shows that dealers might prefer to conceal information about past trades, which allows them to cream-skim the following orders at the expense of other dealers. With the opacity, dealers are willing to make more attractive offers at early stages to obtain this informational advantage.[8] In our paper, by concealing covers, traders might benefit at the expense of dealers. With the opacity, a winning dealer is at an informational disadvantage and, thus, dealers offer less attractive pricing in the early stages.

In the recent models of Kakhbod and Song (2020) and Vairo and Dworczak (2023), informed long-lived dealers trade with myopic uninformed traders who learn from past transactions' data. Interestingly, Kakhbod and Song (2020) show that a dealer finds it easier to sustain opaque pricing, i.e., the one that does not reveal her private information, if traders can observe only past order flow compared to the case in which traders can observe both past order flow and pricing. Vairo and Dworczak (2023) consider a setup with two dealers, who could observe each other's quotes if the market is pre-trade transparent, and argue for the symmetry of information across dealers. Complementary to these papers, we study post-trade transparency choice of long-lived investors trading with dealers. We show that, by concealing covers, investors might endogenously prevent dealers from obtaining symmetric information.

Several empirical studies measure the effects of greater post-trade transparency on the corporate bond market by analyzing the introduction of TRACE that disseminates data on completed trades. Bessembinder et al. (2006); Edwards et al. (2007); Goldstein et al. (2007) find reductions in round-trip trading costs for investors after TRACE starts reporting past transactions. In addition, Asquith, Covert and Pathak (2019) find no change or a decrease in number of trades, with greater effects for high-yield bonds. Lewis and Schwert (2021) document the expansion in dealer networks and decreased risk of market-making.

While we study post-trade investor disclosure, a number of papers in the literature on

[7] See also for Foucault, Pagano and Röell (2013) for a textbook treatment of the setting.

[8] In a setting with no post-trade transparency, Pinter, Wang and Zou (2022) study the trade-off between this information chasing effect and adverse selection.

OTC markets consider voluntary investor disclosure before trade with dealers. Glode, Opp and Zhang (2018) demonstrate that an informed investor might disclose some of her private information before trade to prevent inefficient screening by a counterparty endowed with market power. On the other hand, Baldauf and Mollner (2023) show that agents could limit their pre-trade information revelation by contacting only a few dealers and requesting two-way quotes.

More generally, our paper contributes to the literature on the multi-dealer electronic trading platforms in OTC markets and their design. Hendershott and Madhavan (2015) study the choice between a bilateral negotiation and an RFQ trading mechanism in the corporate bond market, and show that the later is preferred for more active, liquid bonds. Hendershott and Madhavan (2015) also highlight the importance of the sealed-bid nature of the RFQ mechanism that helps prevent information leakage and tacit collusion among dealers. Using more recent data, O'Hara and Zhou (2021) provide a comprehensive analysis of the impact of electronic trading on the corporate bond market. Hendershott et al. (2021) argue that, although electronic platforms in the corporate bond market might have limited ability to spur direct trade between investors, such platforms serve as an important tool for new dealers competing in liquidity provision. Among other papers discussing limits of multi-dealer platforms are Wang (2023) and Yueshen and Zou (2022).

The auctions literature traditionally has focused on information management by the seller without repeated trade. Key works such as Milgrom and Weber (1982b) and Esó and Szentes (2007) have highlighted the general principle that the auctioneer should typically favor full information disclosure ex-ante. By contrast, our paper focuses on ex-post information management. Hausch (1987) provides examples of asymmetric common-value auctions where asymmetry between the bidders mitigates the winner's curse, and the seller may choose not to provide additional information to the less-informed bidder. While in our baseline model the nature of bidder asymmetry is different, we illustrate in an example similar to Hausch (1987) that this incentive for the seller to conceal information may potentially emerge in the dynamic auction setting. A series of papers have also stressed that the auction holder may inject additional uncertainty into the sale process by using a stochastic allocation rule (see Kong, Perrigne and Vuong (2022) and Allen, Clark, Hickman and Richert (2023) for recent examples), finding that, broadly speaking, this additional uncertainty tends to be undesirable for allocative efficiency or seller revenues. In this paper, we highlight the situation where keeping the buyer uncertain benefits the seller.

The balance of the paper is organized in the following way. In the next section, we introduce the model setup. Section 3 presents the equilibrium analysis and derives conditions under which it is optimal to reveal or conceal a covering bid. This section also studies how the choice of transparency affects welfare and market liquidity. Section 4 establishes the conditions under which commitment to ex-post disclosure in the OTC market can be supported through trading relationships between investors and dealers. In Section 5, we extend the model to more general settings and show that our conclusions continue to hold. Section 6 provides a discussion of the results and important assumptions of the model. The last section concludes.

2 Model Setup

We first set up the model and provide its interpretation in the context of the corporate bond market.

Agents and values: There are two agents, a seller and a buyer, who trade in two rounds $t = 1, 2$. Both agents are risk neutral and have a discount factor δ.

The seller offers for sale one bond in each of the two periods. The value of a bond to the seller in period t is v_t. The seller knows the value v_t. One can think of v_t as the best quote that the seller is able to obtain from other buyers or as a fair market value of the bond. In this interpretation, the seller could be an institutional investor who initiated an RFQ or previously contacted other dealers and, through this process, learned their quotes.

The buyer attempts to buy the bonds in each of the two periods. The value of a bond to the buyer in period t is $v_t + B$ where $B > 0$ is a private component of the value to the buyer. The buyer cannot observe v_t but he knows the distribution of v_t. One can think of the buyer as a dealer in the bond market who has to quickly respond to the RFQ initiated by the seller. In this case, the dealer has some information about the common value component v_t but is less informed than the market. The dealers' private value component B could be motivated by the ability of the dealer to resell the bonds among his clients at a markup B.[9]

In the baseline model, we assume that the value of the bonds offered for sale at $t = 1$ and $t = 2$ is the same, i.e., $v_1 = v_2 = v$.[10] The value v is drawn before the first period from a commonly known distribution with the cumulative distribution function (CDF) denoted by $F(v)$ and associated probability density function (PDF) denoted by $f(v)$, which is positive and finite everywhere on its support $[\underline{v}, \bar{v}]$.

Trade: In each of the two periods, the buyer makes a take-it-or-leave-it bid p_t to the seller for the bond being offered for sale in a given period. The seller then decides in the same period whether to accept the bid. That is, as a standard practice in the bond markets, the dealer responds to the RFQ in each period only once and cannot update his quote (see, e.g., Hendershott and Madhavan, 2015), while the seller is committed to choose one of the received quotes — either the dealer's bid or the best bid of other dealers. Thus, agents cannot renegotiate.

Post-trade transparency: We assume that the buyer can observe the seller's value $v_1 = v$ if the buyer's bid p_1 is rejected in the first period. This assumption can be motivated by the post-trade price transparency of the bond market. That is, if the seller rejects the buyer's quote p_1 and picks up the best alternative quote v_1, the trade at the price v_1 is recorded in the TRACE and could be observed by all market participants.

On the other hand, if the seller accepts the buyer's bid p_1 in the first period, the seller's value $v_1 = v$ remains private information of the seller unless there is additional disclosure. The accepted bid allows the buyer to learn something about the seller's value but does not reveal the value fully. Thus, if the dealer wins the RFQ in the first period, the seller's value v_1 is a cover bid and it is private information of the seller unless the seller discloses the cover to the buyer.

We assume that, at $t = 0$, before the value of the bond becomes known to the seller, she commits to either i) disclose the cover, or ii) hide it after the trade in the first period. In

[9] Di Maggio et al. (2017) documents that dealers charge significantly higher markups when trading with their clients as opposed to other dealers.

[10] In Section 5.1, we show that the results are robust if the bond values could change between the two periods. That is, if the seller offers the same bonds in the two periods but their market value changes, or if the bonds offered in the two periods are different but have correlated values.

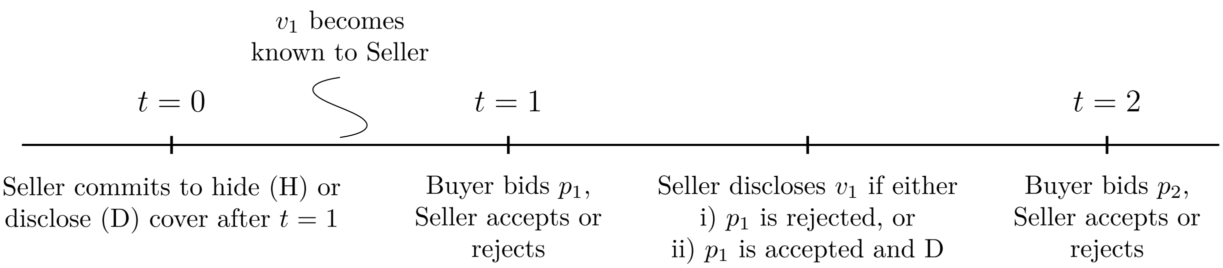

Figure 1: Model timeline.

turn, the buyer knows whether the seller discloses the cover after $t = 1$ before submitting bid p_1 (see the timeline in Figure 1). This can be motivated by the fact that, in the bond market, the identity of a seller initiating an RFQ is usually known to participating dealers and some institutional investors develop reputations for not disclosing covers after trade.[11] In the analysis below, we determine conditions under which the seller prefers to keep the cover hidden.

In line with the literature (e.g. Myerson, 1981), we impose a common regularity condition[12] on the distribution of the bond values that ensures that first-order conditions in the trading game are sufficient conditions for the buyer's optimal bidding decisions.

Assumption 1. *The distribution of the bond values v is such that the ratio*

$$r(v) = \frac{f(v)}{F(v)} \tag{1}$$

is strictly decreasing on the support $[\underline{v}, \bar{v}]$.

Many standard distributions, including uniform and truncated normal, satisfy Assumption 1 and, throughout our analysis, we use such distributions to illustrate our results. Finally, to ensure that the probability of trade in the equilibrium is positive, in the paper, we assume that $\lim_{v \downarrow 0} r(v)$ is sufficiently high.

[11] See Reuters (2018).

[12] In the auctions literature, the analogous assumption is a monotone hazard rate, which is commonly made for buyer valuations in optimal auction design (Myerson, 1981). For symmetric distributions, the condition in Assumption 1 is equivalent to assuming that $(-v)$ satisfies the monotone hazard rate condition.

3 Equilibrium Analysis

To establish the unique subgame-perfect Nash equilibrium of the game, we first derive the equilibrium strategies and profits of the agents in the two cases — with a disclosed cover and with a hidden cover. We then determine the seller's optimal action at $t = 0$ by comparing her expected profits in the two cases.

3.1 Disclosed Cover

In this section, we consider the case in which the seller commits to disclose the cover to the buyer after accepting trade in the first period. In this case, the buyer learns the seller's value v irrespective of the trading outcome in the first period. Indeed, if the seller rejects trade in the first period, v is the execution price which is mandatorily disclosed. If the seller accepts trade in the first period v is the cover.

Knowing the seller's value, the buyer is able to extract all trade surplus in the second period by bidding $p_2^d = v$.[13] Thus, the seller's profit in the second period is equal to $s_2^d(p_2^d) = 0$ while the buyer earns $b_2^d(p_2^d) = B$.

In the first period, if the buyer bids p_1 the seller accepts trade only if it is profitable, i.e., if $v \leq p_1$, because her first-period actions do not affect her second-period profit. Thus, the buyer's total profit evaluated at the first period when bidding p_1 is

$$
\begin{aligned}
b^d(p_1) &= \mathbb{E}[v + B - p_1 | v \leq p_1] \Pr(v \leq p_1) + \delta b_2^d(p_2^d) && (2)\\
&= \int_0^{p_1} v f(v) dv + (B - p_1) F(p_1) + \delta B. && (3)
\end{aligned}
$$

Then, the buyer's marginal benefit of increasing the bid p_1 is

$$
(b^d)'(p_1) = B f(p_1) - F(p_1) = B F(p_1) \left(r(p_1) - \frac{1}{B} \right). \tag{4}
$$

Thus, in the first period, the buyer faces a standard trade-off. On the one hand, by making marginally higher bids, the buyer marginally increases the chance that the seller will

[13]Throughout the paper, we use superscripts $\{d, h\}$ to denote the cases with a disclosed and a hidden cover, respectively.

accept the trade, delivering the buyer additional gains from trade B, which results in a total marginal benefit equal to $Bf(p_1)$. On the other hand, by making marginally higher bids, the buyer pays marginally more to all sellers who accept the trade, which results in a total marginal cost of $F(p_1)$. By Assumption 1, $r(v)$ is strictly decreasing. Therefore, if $r(\bar{v}) = \frac{f(\bar{v})}{F(\bar{v})} > \frac{1}{B}$, the marginal benefit is positive for any p_1 and the buyer optimally bids $p_1^d = \bar{v}$. Alternatively, the buyer's optimal bid is given by the interior solution:

Proposition 1. *If the seller discloses the cover after the first period, the optimal bid of the buyer in the first period p_1^d is given by*

$$Bf(p_1^d) - F(p_1^d) = 0 \tag{5}$$

if $r(\bar{v}) \leq \frac{1}{B}$. Otherwise, the optimal bid is $p_1^d = \bar{v}$. The optimal bid of the buyer in the second period is $p_2^d = v$.

If a cover is revealed, the execution prices, which are disclosed via TRACE, have a declining or flat pattern. If the buyer's bid is accepted in the first period, the seller trades the bond for $p_1^d > v$ at $t = 1$ and for $p_2^d = v$ at $t = 2$. On the other hand, if the buyer's bid is rejected in the first period, the seller trades the bond for v at both $t = 1$ and $t = 2$.

Under the optimal strategies, the seller's total profit evaluated at the first period is

$$s^d(p_1^d) = \mathbb{E}[p_1^d - v | v \leq p_1^d] \Pr(v \leq p_1^d) + \delta s_2^d(p_2^d). \tag{6}$$

3.2 Hidden Cover

Next, we consider the case in which the seller commits to keep the cover hidden from the buyer after accepting trade in the first period. In this case, the buyer learns the seller's value v only if the seller rejects the buyer's bid in the first period. Indeed, if the seller rejects trade in the first period, v is the execution price which is mandatorily disclosed. If the seller accepts trade in the first period v is the cover.

As in the case of the previous section, if the buyer's bid is rejected in the first period, the buyer is able to extract all trade surplus in the second period by quoting $p_2^h = v$, which yields the seller's profit in the second period of $s_2^h(p_2^h) = 0$ and the buyer's profit of $b_2^h(p_2^h) = B$.

Alternatively, if the buyer's bid p_1 is accepted in the first period, the buyer does not learn the exact value of v. However, the buyer infers that the seller's value is $v \in [\underline{v}, \hat{v}]$, where $\hat{v}$ denotes the value of the marginal seller that accepts trade in the first period. The following lemma pins down this marginal seller.

Lemma 1. *The bond value of the marginal seller that accepts trade in the first period is equal to the buyer's bid in the first period, i.e.,* $\hat{v} = p_1$.

Proof. Suppose that $\hat{v} < p_1$. In this case, consider a seller with a value $v \in (\hat{v}, p_1)$. This seller rejects the trade in the first period and, as a result, has to disclose her value to the buyer before the second period. Thus, the seller's profit in both periods is 0. It is clear that the seller could profitably deviate by accepting the buyer's bid p_1 and thereby earning $p_1 - v > 0$ in the first period and a non-negative profit in the second period.

Alternatively, suppose $\hat{v} > p_1$. In this case, consider the marginal seller with the value $\hat{v}$. This seller accepts trade in the first period at a loss, which allows her to keep her value hidden from the buyer in the second period. In the second period, the buyer would optimally choose a bid $p_2 \leq \hat{v}$. Thus, the seller either trades for a profit of zero if $p_2 = \hat{v}$ or rejects the trade otherwise. Overall, the seller's profit in the first period is negative, $p_1 - \hat{v} < 0$, while in the second period it is 0. Therefore, the seller could profitably deviate by rejecting the buyer's bid in the first period and earning a zero total profit over the two periods. As a result, in the equilibrium, $\hat{v} = p_1$. $\square$

By Lemma 1, the seller does not have incentives to strategically reject or accepts the buyer's bid and if the buyer bids p_1 in the first period, the seller accepts trade only when $v \leq p_1$. Thus, in the case with a hidden cover, the buyer's total profit evaluated at the first period when bidding p_1 is

$$b^h(p_1) = \mathbb{E}[v + B + \delta b_2^h(p_2^h(p_1)) - p_1 | v \leq p_1] \Pr(v \leq p_1) + \delta B \Pr(v > p_1) \tag{7}$$

$$= \mathbb{E}[v + B + \delta(b_2^h(p_2^h(p_1)) - B) - p_1 | v \leq p_1] \Pr(v \leq p_1) + \delta B, \tag{8}$$

where $b_2^h(p_2)$ denotes the buyer's profit in the second period after the seller accepts the bid p_1 in the first period and $p_2^h(p_1)$ denotes the optimal second-period bid as a function of the accepted bid p_1.

The total profit in (8) is almost the same as the buyer's total profit in the case with a disclosed cover in (2) apart from the term $\mathbb{E}[\delta(b_2^h(p_2^h(p_1)) - B)|v \leq p_1] \Pr(v \leq p_1)$. This term is negative as $b_2^h(p_2) \leq B$ for any p_2 because the buyer's profit is no greater than the total gain from trade in the second period. Thus, in case with a hidden cover, the buyer makes weakly lower — more aggressive — bids in the first period compared to the case with a disclosed cover, i.e., $p_1^h \leq p_1^d$.

Intuitively, by winning trade in the first period, compared to losing, the buyer earns lower profit in the second period. This means that the buyer has lower incentives to win trade in the first period and, as a result, submits lower bids. This is in contrast to the case with a disclosed cover in which the buyer does not become disadvantaged from winning and earns the same profit in the second period irrespective of the outcome in the first period.

Next, we find the buyer's optimal bid $p_2^h(p_1)$ and profit $b_2^h(p_2^h(p_1))$ in the second period after the seller accepts the buyer's bid p_1 in the first period. By Lemma 1, the buyer rationally infers that the seller's value v^a is distributed on $[\underline{v}, p_1]$ and the density of this truncated distribution is $\frac{f(v)}{F(p_1)}$. Because it is the last period of the game, when the buyer bids p_2, the seller accepts trade only if it is profitable, i.e., if $v^a \leq p_2$. Thus, the buyer's profit when bidding p_2 is

$$
\begin{aligned}
b_2^h(p_2) &= \mathbb{E}[v^a + B - p_2|v^a \leq p_2] \Pr(v^a \leq p_2) && (9)\\
&= \left(\int_0^{p_2} vf(v)dv + (B - p_2)F(p_2)\right) / F(p_1). && (10)
\end{aligned}
$$

Then, the buyer's marginal benefit of increasing the bid p_2 is

$$
(b_2^h)'(p_2) = \frac{Bf(p_2) - F(p_2)}{F(p_1)}. \tag{11}
$$

This benefit is proportional to the marginal benefit for the first period trade in the case with a disclosed cover (4), which determines the optimal bid p_1^d. Thus, for any buyer's bid $p_1 \leq p_1^d$, by the monotonicity in the Assumption 1, we have $Bf(p_1) - F(p_1) \geq 0$. Therefore, the marginal benefit (11) is weakly positive at the upper boundary of the truncated distribution,

p_1, and the buyer optimally bids $p_2^h(p_1) = p_1$. Intuitively, if the buyer makes a more aggressive bid p_1 in the case with a hidden cover, compared to p_1^d, which is the optimal bid in a one-shot game, and p_1 is accepted, the buyer does not have incentives to lower the bid p_2 in the remaining one-shot game of the second period.

Consequently, if we make a conjecture, which is verified below, that $p_1 \leq p_1^d$, the buyer's profit in the second period is

$$b_2^h(p_2^h(p_1)) = \int_0^{p_1} v \frac{f(v)}{F(p_1)} dv + (B - p_1). \tag{12}$$

Plugging this into the buyer's total profit evaluated at the first period (8), yields

$$\begin{aligned} b^h(p_1) &= \int_0^{p_1} (v + B + \delta[b_2^h(p_2^h(p_1)) - B] - p_1) f(v) dv + \delta B & (13) \\ &= (1+\delta) \int_0^{p_1} v f(v) dv + [B - (1+\delta)p_1] F(p_1) + \delta B. & (14) \end{aligned}$$

Then, the buyer's marginal benefit of increasing the bid p_1 is

$$(b^h)'(p_1) = Bf(p_1) - (1+\delta)F(p_1) = BF(p_1)\left(r(p_1) - \frac{1+\delta}{B}\right). \tag{15}$$

This marginal benefit is similar to the one for the first period trade in the case with disclosure (4) except there is now an additional marginal cost of $\delta F(p_1)$. Note that there is no additional marginal benefit due to gains from trade because the buyer is able to capture the gains in the second period even if the bid p_1 is rejected. Because $\delta > 0$, the greater cost implies that the benefit of marginally raising price p_1 is smaller in the case with a hidden cover. Therefore, the buyer makes weakly more aggressive bids in the first period, i.e, $p_1^h \leq p_1^d$. Thus, the above conjecture is verified and we can write down the buyer's optimal strategy (see Figure 2).

Proposition 2. *If the seller keeps the cover hidden after the first period, the optimal bid of the buyer in the first period p_1^h is given by*

$$Bf(p_1^h) - (1+\delta)F(p_1^h) = 0 \tag{16}$$

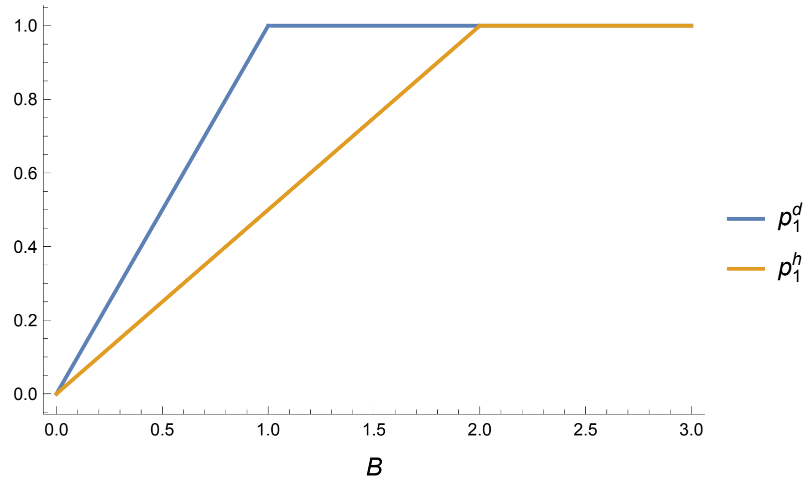

Figure 2: The buyer's optimal bids in the first period in the cases with a disclosed cover, p_1^d, and a hidden cover, p_1^h. In the graph, we consider a setting with $v \sim U[0,1]$ and $\delta = 1$.

if $r(\bar{v}) \leq \frac{1+\delta}{B}$. *Otherwise, the optimal bid is* $p_1^h = \bar{v}$. *The optimal bid of the buyer in the second period is* $p_2^h = v$ *if the trade is rejected in the first period and it is* $p_2^h = p_1^h$ *if the trade is accepted in the first period.*

Interestingly, in the equilibrium, if the buyer's bid p_1^h is accepted in the first period and, thus, he learns that the seller's value $v^a \in [\underline{v}, p_1^h]$, the buyer does not find it optimal to decrease his bid in the second period by setting p_2^h lower than p_1^h. That is, the buyer does not learn additional information about the seller's value after the second period. Indeed, the buyer has lower incentives to learn at $t = 2$ because he can afford more aggressive experimentation with lower prices in the first period — if the trade is rejected at $t = 1$ the buyer gets to learn the exact value of v.

If a cover is hidden, the execution prices, which are disclosed via TRACE, have a flat pattern. If the buyer's bid is accepted in the first period, the seller trades the bond for $p_1^h > v$ at $t = 1$ and for $p_2^h = p_1^h$ at $t = 2$. On the other hand, if the buyer's bid is rejected in the first period, the seller trades the bond for v at both $t = 1$ and $t = 2$.

Under the optimal strategies, the seller's total profit evaluated at the first period is

$$s^h(p_1^h) = (1+\delta)\,\mathbb{E}[p_1^h - v | v \leq p_1^h]\,\Pr(v \leq p_1^h). \tag{17}$$

In the next section, we compare this profit to the seller's profit in the case with a disclosed cover $s^d(p_1^d)$ which is given by (6).

3.3 Seller's Choice of Transparency

In this section, we determine whether, at $t = 0$, the seller prefers to commit to disclose or keep a cover hidden after the trade in the first period.

When choosing between the two cases, the seller faces the following trade-off. On the one hand, if the seller decides to commit to hide a cover after the first period, she is able to retain some private information if the trade is accepted in the first period, which allows her to earn a higher profit in the second period. On the other hand, such commitment implies that the buyer makes more aggressive — lower bids in the first period, which leads to a lower seller's profit in the first period.

Formally, to find the optimal decision of the seller, we need to compare her expected profits when a cover is disclosed (6):

$$s^d(p_1^d) = \int_0^{p_1^d} (p_1^d - v) f(v) dv + \delta \cdot 0 \tag{18}$$

and when it is hidden (17):

$$s^h(p_1^h) = \int_0^{p_1^h} (p_1^h - v) f(v) dv + \delta \int_0^{p_1^h} (p_1^h - v) f(v) dv. \tag{19}$$

In particular, if the seller chooses the latter case over the former, the difference in her profit of the first period is $\Delta_1 = \int_0^{p_1^h} (p_1^h - v) f(v) dv - \int_0^{p_1^d} (p_1^d - v) f(v) dv \leq 0$, which is negative due to the more aggressive bidding by the buyer $p_1^h \leq p_1^d$. However, the difference in her profit of the second period is $\Delta_2 = \delta \int_0^{p_1^h} (p_1^h - v) f(v) dv - \delta \cdot 0 \geq 0$, which is positive due to the retained private information. The total difference in profits is $\Delta_s = \Delta_1 + \Delta_2$ and the seller optimally chooses to hide a cover whenever $\Delta_s \geq 0$.

Thus, if the buyer's optimal bid does not decrease significantly when the cover information is not available, i.e., if p_1^h and p_1^d are relatively close to each other, Δ_1 is close to zero. At the same time, Δ_2 is positive and bounded away from zero because it depends only on p_1^h. In such cases, the total difference in the profit $\Delta_s \geq 0$ and the seller prefers to hide the cover.

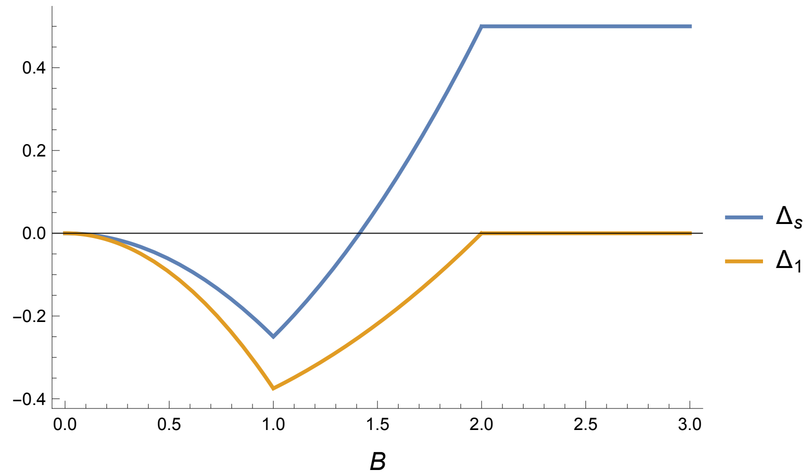

Figure 3: The difference in the seller's total profit, Δ_s and the seller's first-period profit, Δ_1, between the cases with a hidden cover and a disclosed cover. In the graph, we consider a setting with $v \sim U[0,1]$ and $\delta = 1$.

The next two propositions show that this is the case if the private benefit B is high enough or if the distribution of v is such that $r(v)$ is sufficiently steep at the right tail, i.e., if v is relatively less risky.

Proposition 3. *There is $\hat{B}$ such that, for any private benefit $B \geq \hat{B}$, the seller finds it optimal to commit at $t = 0$ to keep the cover hidden after $t = 1$, i.e., $\Delta_s \geq 0$ for such B.*

Intuitively, if the buyer's private benefit is very large, he would be extremely reluctant to risk the possibility of being rejected by the seller. Therefore, even if no information is revealed after the trade, the buyer will try to make sure the bid is good enough to be accepted, which improves the seller's expected payoff. In the limit, for very high B, the buyer makes the least aggressive bids in both cases, with a disclosed cover and with a hidden cover, i.e., $p_1^h = p_1^d = \bar{v}$. In this scenario, $\Delta_1 = 0$ while $\Delta_2 = \delta(\bar{v} - \mathbb{E}\, v) > 0$ and, thus, $\Delta_s > 0$, which means that the seller keeps the cover hidden. On the other hand, when the gains from trade are small, compared to the case of disclosed covers, the buyer would significantly discount his offer. This is costly from the seller's perspective and she would rather commit to disclosing the cover. Figures 3 and 4 provide an illustration for this result.

Proposition 4. *If the distribution of the bond values v is such that $r(v)$ is sufficiently steep at the right tail, i.e., if v is relatively less risky, the seller finds it optimal to commit at $t = 0$ to keep the cover hidden after $t = 1$, i.e., $\Delta_s \geq 0$ for such distributions.*

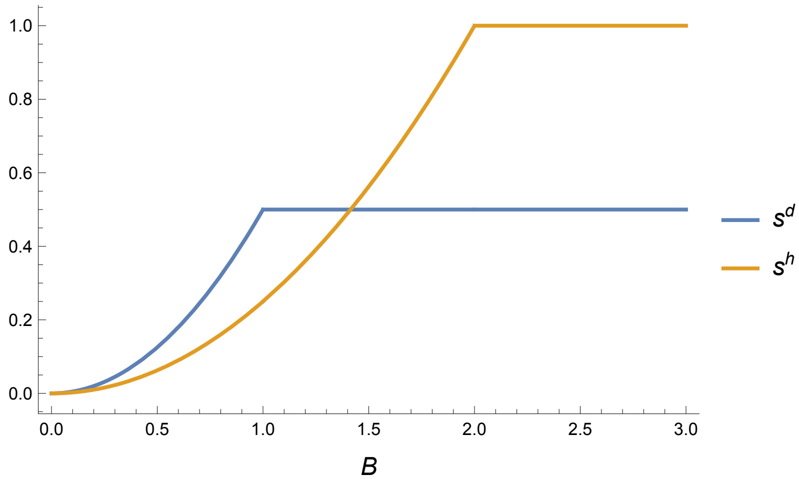

Figure 4: The seller's total profit in the cases with a disclosed cover, s^d, and a hidden cover, s^h. In the graph, we consider a setting with $v \sim U[0,1]$ and $\delta = 1$.

Intuitively, if $r(v)$ is sufficiently steep at the right tail, the distribution of the bond value v has a thin right tail, i.e., the bond values have lower variation. Therefore, the value of additional information contained in the cover is relatively small to the buyer. This implies that the buyer would not significantly discount his bid in the first period if the cover information is not provided. Indeed, for such distributions, the buyer's optimal bids p_1^d and p_1^h, which are given by (5) and (16), are relatively close to each other. Thus, $\Delta_s \geq 0$ and the seller chooses to hide the cover. Figure 5 provides an illustration for the this result. In the figure, we consider a symmetric distribution on $[0,1]$ for which parameter a controls the steepness of $r(v)$ at the distribution's tails.[14] Higher a indicates thinner tails.

Both Propositions 3 and 4 show that there are cases in which the seller prefers to keep the cover information private even though the buyer is willing to quote a weakly higher bid if this information is disclosed after the trade in the first period. In other words, in these cases, the seller's value of retaining the information is higher than what the buyer is willing to pay to obtain this information.

[14]The density of the distribution is such that $f(\frac{1}{2}) = 2 - \frac{1}{a}$ and $f(1) = \frac{1}{a}$, so, for higher a, it is more concentrated around its mean $\frac{1}{2}$. $r(v) = \frac{(3-4a)+4v(a-1)}{(a-1)+v(3-4a)+2v^2(a-1)}$ for any $v \geq \frac{1}{2}$.

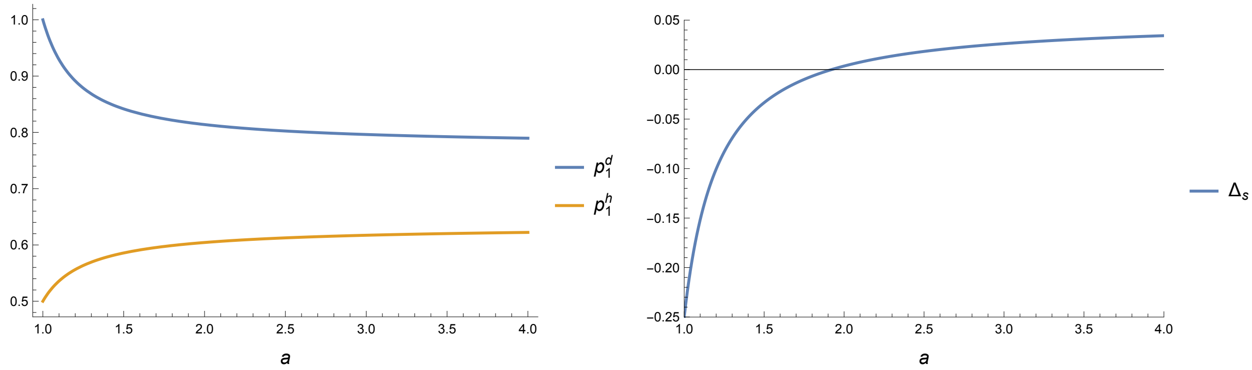

Figure 5: Left panel: The buyer's bids in the first period in the cases with a disclosed cover, p_1^d, and a hidden cover, p_1^h. Right panel: The seller's total profit improvement Δ_s for the case with a hidden cover compared to the case with a disclosed cover. In the graph, we consider a setting with $f(v) = \left[\frac{1}{a} + v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(0 \leq v \leq \frac{1}{2}) + \left[\left(4 - \frac{3}{a}\right) - v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(\frac{1}{2} < v \leq 1)$, $B = 1$, and $\delta = 1$. $a = 1$ corresponds to the uniform $[0, 1]$ distribution and higher a indicates thinner tails.

3.4 Welfare

In this section, we study whether it is more efficient to disclose or hide the cover from the social perspective. In the case with a disclosed cover, the total surplus, which is a sum of the buyer's and seller's profits, is

$$W^d(p_1^d) = b^d(p_1^d) + s^d(p_1^d) = BF(p_1^d) + \delta B. \tag{20}$$

In this case, the trade could break down only in the first period while in the second period trade always happens. Indeed, because the buyer learns the seller's exact value v after the first period, his bid in the second period $p_2^d = v$ is always accepted.

In the case with a hidden cover, the total surplus is

$$W^h(p_1^h) = b^h(p_1^h) + s^h(p_1^h) = BF(p_1^h) + \delta B. \tag{21}$$

Interestingly, as in the former case, here too, the trade could break down only in the first period. Specifically, if the seller accepts trade in the first period, the buyer makes the same bid in the second period since $p_2^h(p_1^h) = p_1^h$ and, by Lemma 1, this implies that the seller accepts trade also in the second period. Alternatively, if the trade is rejected in the first period, the buyer learns the seller's exact value v and his bid $p_2^d = v$ is accepted.

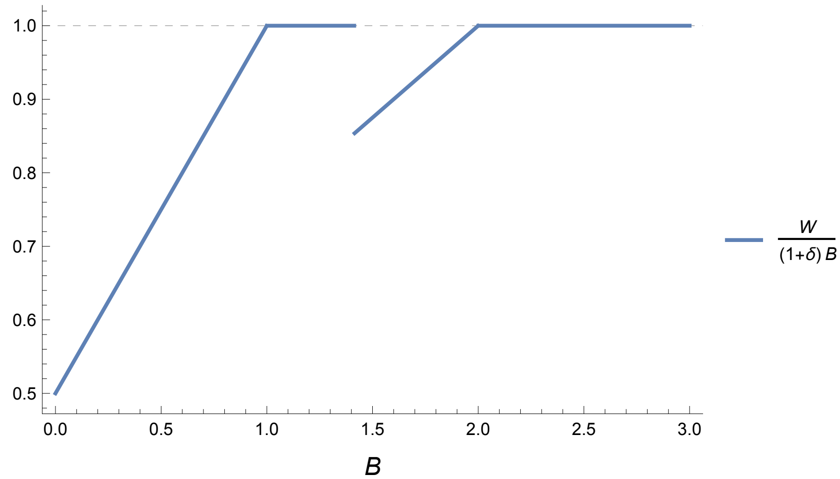

Figure 6: The ratio of the total surplus W to the maximum gain from trade $(1+\delta)B$ under the seller's optimal choice between a disclosed cover and a hidden cover. In the graph, we consider a setting with $v \sim U[0,1]$ and $\delta = 1$.

Therefore, the difference in the total surplus between the two cases is only due to the difference in the probability of trade in the first period. Because, the buyer is weakly more aggressive in the latter case, i.e., $p_1^h \leq p_1^d$. The trade is weakly more likely to break down. This implies that the total surplus with a hidden cover is weakly smaller than that in the case with a disclosed cover. Thus, disclosure is socially efficient even though it could reduce the seller's profit.

By Proposition 3, the seller prefers to hide a cover for any $B > \hat{B}$ where $\hat{B}$ is such that the seller's profits in the cases with a disclosed and with a hidden cover are equal, $s^d(p_1^d) = s^h(p_1^h)$. Because $p_1^d > p_1^h$ at this point, the social surplus drops discontinuously at $\hat{B}$ if the seller switches her choice of transparency at $t = 0$ from the regime with a disclosed cover to the regime with a hidden cover (see Figure 6). This decrease in the social surplus is equal to the decrease in the buyer's profit.

Proposition 5. *i) The social surplus is higher in the case with a disclosed cover compared to the social surplus in the case with a hidden cover if the buyer's first-period bid is more aggressive in the latter case, i.e., if* $p_1^h < p_1^d$. *ii) Unlike for a one-period trade, higher potential gains from trade, i.e., higher* B, *do not necessarily imply higher total surplus.*

The above results imply that disclosure of cover information is an important determinant of market liquidity that affects social surplus. Crucially, under conditions identified in

Propositions 3 and 4, the seller might privately prefer to keep the cover hidden which harms realization of gains from trade.

4 Commitment to Ex-Post Disclosure

The results of the previous section are established under the assumption that the seller is committed to the information disclosure rules chosen before the trade. While digital platforms, where disclosure is encoded and automatic, naturally provide such assurance, OTC markets do not have built-in commitment. In Section 4.1, we verify that the buyer and the seller have sufficiently misaligned incentives, so that without commitment the seller finds it optimal to deviate from the promised disclosure ex-post. In Section 4.2, we extend our baseline model to allow repeated interactions and show that it is possible to sustain commitment in the OTC market only if the trading relationship between the agents is sufficiently strong. Therefore, sellers trading infrequently, with weak relationships in the OTC market, cannot commit to share information ex-post even through these sellers are precisely the ones who would benefit most from such commitment.

4.1 Commitment in One-Shot Game

First, we illustrate the problem of commitment in the OTC markets using our baseline model. Consider a seller who promised to reveal the cover to the buyer after the first trade. If the buyer takes the information received from the seller at face value, the seller may potentially be interested in misreporting the truth. Indeed, if the value v of the bond is below the original bid p_1 offered by the buyer, the seller benefits from reporting the bond value as p_1. In the second period of trade, the buyer will make the same offer provided he believes the disclosure. This offer yields a positive profit for the seller which is greater than her profit of 0 in case if she reports truthfully.

4.2 Commitment in Repeated Game

One natural countervailing force to the incentive to misreport is generated in repeated interactions: if the buyer later discovers that the value has been inflated, he may find a way to reduce the seller's payoff in subsequent trades. Repeated interactions, however, do not necessarily fully offset this incentive for the players who are too impatient. Below, we confirm that coordination breakdown is possible in our model: for sellers with sufficiently small discount factors, the incentive to misreport can outweigh the incentive to maintain a relationship with the buyer.

Recall the inequality determining whether the seller prefers cover disclosure:

$$\int_0^{p_1^d} (p_1^d - v) f(v) dv \geq (1+\delta) \int_0^{p_1^h} (p_1^h - v) f(v) dv \tag{22}$$

Here, we are interested whether this inequality holds for small discounts δ. It is not immediately clear that it is satisfied when δ is close to 0. On the one hand, the seller's future profit is discounted, making disclosure more attractive. On the other hand, the gap between p_1^d and p_1^h is also decreasing as δ gets smaller. In the limit where $\delta = 0$, we have $p_1^d = p_1^h$, and the seller is indifferent between her options. However, if the seller and the buyer have distinct discounts, we can make a more precise prediction.

Suppose now that the seller's discount over time is δ_s, and the buyer's discount is δ_b. In the corporate bond setting, such a distinction is natural, since these two types of agents play different roles in the market: one side initiates trades while the other facilitates them. With this notation, we have $p_1^h = p_1^h(\delta_b)$, $p_1^d = p_1^d(\delta_b)$, and in equation (22) the right-hand side discount is replaced with δ_s. Note that, conditional on δ_b, we have $p_1^d > p_1^h$. Therefore, the seller prefers disclosure for all small enough δ_s.[15]

To formalize the argument that in repeated interactions, sellers with low δ_s may fail to commit, we set up a game where the same buyer and seller trade pairs of bonds in one-shot games analogous to the baseline model we study in the previous sections. Let $t = 0, 1, \ldots$

[15] When $\delta_s = \delta_b$, the inequality can still remain true for small discount values. One such example in the Appendix is the uniform distribution. It is possible, however, to construct distributions for v such that the inequality fails and the seller prefers to hide the cover.

index time periods. Each time period, a bond is sold either to the fringe or to the dealer. At even times $2t$, the seller learns the bond valuation v_{2t}, while at odd times $2t+1$, the valuation is inherited from the previous time period: $v_{2t+1} = v_{2t}$. The buyer's valuation $v_{b,t} = v_t + B$ is composed of v_t and the private component B which, for simplicity, is assumed to be constant and known. The dealer who successfully purchases bonds at $2t$ and $2t+1$ learns their exact valuations only at $2t+2$.[16] At the beginning of each time period, the seller can send a message $\hat{v}_t$ to the buyer. She always has the outside option of realizing v_t but she can also accept bid b_t submitted by the buyer during trade period t, which is a take-it-or-leave-it offer from the buyer. Each transaction price at time t becomes public information at $t+1$. Across pairs of trades, bond values v_{2t} are i.i.d. continuously distributed with finite support on $[\underline{v}, \overline{v}]$. The seller's time discount is δ_s, and the buyer's time discount is δ_b.

Repeated games with learning may have many subgame-perfect Nash equilibria. To discuss commitment, we restrict our attention to a subset of strategies where the seller can choose what she reports to the buyer, while the buyer plays a grim-trigger strategy. Specifically, the buyer acts as in the one-shot Nash equilibrium found in our baseline model for as long as the seller reports truthfully; if at some point in time $2t+2$ the buyer finds out the previous message $\hat{v}_{2t+1}$ was untruthful, she stops the trade altogether.

Consider an equilibrium where the seller is trying to replicate the outcomes of the one-shot baseline and report truthfully. Pick some odd time period $2t+1$. The seller's payoff from cooperation is

$$s_{coop} = 0 + \delta_s s^d + \delta_s^3 s^d + \ldots = \delta_s \frac{1}{1-\delta_s^2} s^d \tag{23}$$

where the contemporaneous payoff 0 is due to the seller reporting $\hat{v}_{2t+1} = v_{2t+1}$. However, if $p_1^d > v_{2t+1}$, the seller may misreport in the current period and forfeit all future profits after the buyer finds out about the misreporting. The deviation payoff is equal to

$$s_{deviate} = p_1^d - v_{2t+1} \tag{24}$$

and is achieved by setting $\hat{v}_{2t+1} = p_1^d$. Note now that $s_{coop} \to 0$ when $\delta_s \to 0$, while $s_{deviate}$

[16]The idea of learning by the buyer is that if he never learns bond values, it is unclear what it means for the buyer to have any valuation for a bond. The exact timing assumptions for this learning can be relaxed.

does not depend on δ_s. Therefore, for small enough δ_s the incentives to deviate are stronger. We summarize the implications of this argument in the proposition below.

Proposition 6. *For low enough discount value of δ_s, there is no subgame-perfect Nash equilibrium in which the seller always discloses the cover truthfully.*

Let us now consider impatient sellers with low δ_s. Based on the analysis in Section 3, the quotes p_1^d and p_1^h are such that $p_1^d - p_1^h$ is a positive number that does not depend on δ_s. Consequently, comparing equations (18) and (19) shows that for small enough δ_s, the seller's ex-ante profit from disclosing the cover is higher than from hiding it. Together with the proposition established above, these results imply that OTC markets make it harder for the infrequently trading sellers to benefit from promising to share price information after the sale. The programmatic commitment afforded by the electronic trading venues can potentially benefit these types of traders, which, per analysis in Section 3, could also improve the total welfare in the market.[17]

5 Extensions

5.1 Value Change between $t = 1$ and $t = 2$

In this section, we show that the results of our baseline model are robust if we allow a change in the bond values between $t = 1$ and $t = 2$.

Specifically, we assume that, with probability α, $v_2 = v_1$ while, with complementary probability $1 - \alpha$, v_2 is independently redrawn from the same distribution as v_1. Thus, the baseline model is the special case with $\alpha = 1$. Additionally, we assume that both the seller and the buyer know before $t = 2$ whether the value changed. However, as for v_1 in the first period, we assume that, at $t = 2$, the seller knows v_2 while the buyer only knows its distribution.

Below we analyze how our results change in the extended setup for the two main cases if the seller discloses a cover and if she hides it. However, first, we consider the agent's optimal

[17] Our results do not necessitate that the buyer and the seller have independent discount factors. For instance, the uniform distribution example in Appendix A.3 shows that when the dealer's private value is low enough, all sellers prefer committing to disclosure.

strategies in the second period if $v_2 \neq v_1$. If $v_2 \neq v_1$, any information that the buyer learned in the first period about v_1 becomes irrelevant at $t = 2$. In this case, the buyer bids p_2 knowing only the distribution of v_2 and the seller accepts trade only if it is profitable, i.e., if $v_2 \leq p_2$. Thus, if $v_2 \neq v_1$, the buyer's profit in the second period when bidding p_2 is

$$\mathbb{E}[v_2 + B - p_2 | v_2 \leq p_2] \Pr(v_2 \leq p_2) = \int_0^{p_2} v f(v) dv + (B - p_2) F(p_2). \tag{25}$$

The buyer's marginal benefit of increasing price p_2 is the same as in (4). Therefore, the optimal bid p_2^n is given by (5). For this optimal bid, denote the buyer's and seller's profits by b_2^n and s_2^n, respectively.

5.1.1 Disclosed Cover

If the seller commits to disclose the cover to the buyer after accepting trade in the first period, the buyer learns the seller's value v_1 irrespective of the trading outcome in the first period. With probability α, the buyer uses this information in the same way as in the baseline model while, with probability $1-\alpha$, this information becomes stale. Thus, the buyer's profit in the second period is $b_2^d(p_2^d) = \alpha B + (1-\alpha) b_2^n$ while the seller earns $s_2^d(p_2^d) = (1-\alpha) s_2^n$.

As in the baseline model, the outcome of the trade in the first period, does not affect the buyer's profit in the second period. Thus, in the first period, the buyer bids the same p_1^d, which is given by (5). Overall, if seller discloses the cover, the possibility of the change in the bond values affects only the agent's profits in the second period.

Under the optimal strategies, the seller's total profit evaluated at the first period is

$$s^d(p_1^d) = \mathbb{E}[p_1^d - v | v \leq p_1^d] \Pr(v \leq p_1^d) + \delta(1-\alpha) s_2^n. \tag{26}$$

5.1.2 Hidden Cover

If the seller commits to keep the cover hidden after accepting trade in the first period, the buyer learns the seller's value v_1 only if the seller rejects the buyer's bid in the first period. On the other hand, if the buyer's bid is accepted in the first period, the buyer infers that the seller's value is $v_1 \in [\underline{v}, \hat{v}]$. With probability α, the buyer uses the information learned in the first period in the same way as in the baseline model while, with probability $1-\alpha$, this information becomes stale.

Thus, if the buyer's bid is rejected in the first period, the buyer's profit in the second period is $b_2^d(p_2^d) = \alpha B + (1-\alpha)b_2^n$ while the seller earns $s_2^d(p_2^d) = (1-\alpha)s_2^n$. Alternatively, if the buyer's bid p_1 is accepted in the first period, the buyer's profit in the second period is $\alpha b_2^h(p_2^h(p_1)) + (1-\alpha)b_2^n$, where $b_2^h(p_2)$ denotes the buyer's profit in the second period after the seller accepts the bid p_1 in the first period and $p_2^h(p_1)$ denotes the optimal second-period bid as a function of the accepted bid p_1.

It can be shown that the result of Lemma 1 applies in the extended setup. Therefore, in the case with a hidden cover, the buyer's total profit evaluated at the first period when bidding p_1 is

$$\begin{aligned} b^h(p_1) &= \mathbb{E}[v + B + \delta[\alpha b_2^h(p_2^h(p_1)) + (1-\alpha)b_2^n] - p_1 | v \leq p_1] \Pr(v \leq p_1) \\ &\qquad\qquad\qquad\qquad\qquad\qquad + \delta(\alpha B + (1-\alpha)b_2^n) \Pr(v > p_1^h) \\ &= \mathbb{E}[v + B + \delta\alpha[b_2^h(p_2^h(p_1)) - B] - p_1 | v \leq p_1] \Pr(v \leq p_1) + \delta(\alpha B + (1-\alpha)b_2^n). \end{aligned} \tag{27}$$

As in the baseline model, by winning trade in the first period, compared to losing, the buyer earns lower profit in the second period. This implies that the buyer have lower incentives to win trade in the first period and, as a result, submits lower bids.

As we show below, if $v_2 = v_1$, the buyer's optimal bid $p_2^h(p_1)$ and profit $b_2^h(p_2^h(p_1))$ in the second period after the seller accepts the buyer's bid p_1 are the same as in the baseline model. Thus, the buyer's total profit evaluated at the first period (27) can be rewritten as

$$b^h(p_1) = \int_0^{p_1} (v + B + \delta\alpha[b_2^h(p_2^h(p_1)) - B] - p_1) f(v) dv + \delta(\alpha B + (1-\alpha)b_2^n) \tag{28}$$

$$= (1+\delta\alpha) \int_0^{p_1} v f(v) dv + [B - (1+\delta\alpha)p_1] F(p_1) + \delta(\alpha B + (1-\alpha)b_2^n). \tag{29}$$

Finally, the buyer's marginal benefit of increasing the bid p_1 is

$$(b^h)'(p_1) = Bf(p_1) - (1+\delta\alpha)F(p_1). \tag{30}$$

This marginal benefit is similar to the one in the main model (15) except the marginal cost $(1+\delta\alpha)F(p_1)$ is smaller because $\alpha < 1$. The smaller cost implies that the benefit of marginally raising price p_1 is higher. Therefore, the buyer makes weakly less aggressive bids in the first period if there is a possibility of the change in bond values between the two periods, i.e, p_1^h is higher than in the baseline.

Proposition 7. *If the seller keeps the cover hidden after the first period, the optimal bid of the buyer in the first period p_1^h is given by*

$$Bf(p_1^h) - (1+\delta\alpha)F(p_1^h) = 0. \tag{31}$$

if $r(\bar{v}) \leq \frac{1+\delta\alpha}{B}$. Otherwise, the optimal bid is $p_1^h = \bar{v}$. If $v_2 = v_1$, the optimal bid of the buyer in the second period is $p_2^h = v$ if the trade is rejected in the first period and it is $p_2^h = p_1^h$ if the trade is accepted in the first period. If $v_2 \neq v_1$, the optimal bid of the buyer in the second period p_2^n is given by (5).

Under the optimal strategies, the seller's total profit evaluated at the first period is

$$s^h(p_1^h) = (1+\delta\alpha)\,\mathbb{E}[p_1^h - v|v \leq p_1^h]\Pr(v \leq p_1^h) + \delta(1-\alpha)s_2^n. \tag{32}$$

5.1.3 Seller's Choice of Transparency

The buyer's and seller's strategies in the extended model are very similar to those in the baseline model. Thus, the seller chooses to hide the cover under the similar conditions to those identified in Propositions 3 and 4.

Compared to the baseline, the possibility of the change in the bond values, i.e., the decrease in α, has two opposite effects on the relative value of the seller's profits when the cover is disclosed (26) and when it is hidden (32). On the one hand, lower α implies that the cover information is not as valuable for the buyer. This brings p_1^h closer to p_1^d and, thus, the seller does not lose as much in the first period by hiding a cover. On the other hand, lower α implies lower probability of states when the seller can benefit from the retained information in the second period, which decreases the relative benefit of hiding the cover.

5.2 T Periods of Trade

In this section, we discuss how our results change if there are more than two periods of trade. Specifically, we assume that the agents trade under the assumptions of the baseline model for $T > 2$ periods.

In the case with a disclosed cover, the equilibrium is the same as the one with $T = 2$. In the first period, the buyer bids p_1^d, which is given by (5). Then, he learns the exact value of the bond v because the seller discloses v irrespective of the trading outcome. If the trade is accepted v is disclosed as a cover, and if the trade is rejected v is disclosed as an execution price. Thus, the buyer bids $p_t^d = v$ in all remaining periods $t \geq 2$.

In the case with a hidden cover, the equilibrium is similar to the one with $T = 2$ apart from the change in the buyer's optimal bid in the first period.

Proposition 8. *With T trading periods, if the seller keeps the cover hidden after trade in every period, the optimal bid of the buyer in the first period p_1^h is given by*

$$Bf(p_1^h) - \left(\frac{1-\delta^T}{1-\delta}\right) F(p_1^h) = 0 \tag{33}$$

if $r(\bar{v}) \leq \left(\frac{1-\delta^T}{1-\delta}\right)\frac{1}{B}$. Otherwise, the optimal bid is $p_1^h = \bar{v}$. The optimal bid of the buyer in all subsequent periods $t \geq 2$ is $p_t^h = v$ if the trade is rejected in the first period, and it is $p_t^h = p_1^h$ if the trade is accepted in the first period.

As in the baseline model, the buyer attempts to learn the seller's value only in the first period. That is, the buyer does not decrease his bid p_1^h in all subsequent periods if it was accepted in the first period. Intuitively, the buyer can afford more aggressive experimentation with lower prices in the earlier periods because, if the trade is rejected at period t, the buyer gets to learn the exact value of v and extracts full gains from trade in all remaining $T - t$ periods. Thus, the buyer's incentives to set lower bids are decreasing with the number of remaining periods $T - t$, which means that the buyer sets a very low bid in the first period and does not update it after. The buyer's optimal bid in the first period p_1^h is more aggressive for larger T because incentives to lose increase with T. The only reason why the bid p_1^h is positive is the gain from trade that the buyer tries to capture in the first period.

When choosing whether to reveal or conceal the cover, with $T > 2$, the seller faces the same trade-off as in the baseline model. If the cover is hidden, higher T allows the seller to capture information rents in more rounds of trade. However, these rents decrease for higher T because the buyer submits more aggressive bid in the first period.

6 Discussion

In this section, we discuss some important assumptions behind the model and the model's implications.

Market Fragmentation and Platform Design. Our results highlight that the different choices of investors towards cover disclosure could explain investors' trading venue choices and, thus, be one of the reasons behind market fragmentation. Indeed, if investors would like to credibly disclose covers ex-post they might choose to trade via an electronic RFQ platform, such as MarketAxess or Tradeweb, where disclosure is encoded and automatic. While if investors would like to conceal covers they might choose to trade offline via phone. Thus, investors' choice of transparency might be one of the factors explaining adoption of the RFQ platforms.

From a platform's perspective, offering more flexible trading mechanisms to its investor clients might not be a straightforward solution to the adoption problem. Indeed, as our model indicates, the ex-post disclosure of covers affects the distribution of surplus between investors and dealers. Therefore, offering more flexible mechanisms to investors is likely to affect participation by the dealers. We leave comprehensive modeling of a platform's decision in light of this trade-off for future work.

Multiple Buyers. In our setup, only one buyer is modeled explicitly. This approach permits tractable analysis of repeated trades and ensures that the equilibrium in the trade game is unique. On the other hand, we cannot directly study the strategic incentives of other buyers in a typical corporate bond auction. We argue, however, that the seller may still have the incentive to conceal the cover and, therefore, reduce the appeal of the first auction even in a setting with multiple strategic buyers.

To highlight this possibility, in Appendix A.2, we discuss a setup with two buyers and two auctions. In the first auction, the buyers receive noisy signals about the value of the bond, which is common to both of them, and trade on that private information. In the follow-up auction, the trade repeats one more time after the buyers update their beliefs based on the outcome of the first auction. The difficulty lies in analyzing the second auction, in which the buyers are potentially asymmetrically informed. We rely on the model of Abraham et al. (2020) to establish an equilibrium in this follow-up auction and calculate the seller's expected revenue. We argue that when bidders do not acquire any additional information after the first auction, it is always better for the seller to reveal the covers. However, if the bidders receive additional independent signals before the second trade, it may be in the seller's interest to keep them asymmetrically informed by hiding the runner-up bid.

7 Conclusion

This paper develops a model with repeated asset sales to illustrate investors' incentives for information disclosure in the corporate bond market. In this market, regulation ensures that the accepted offer is revealed for every sale; however, we argue that the offers which are received but not accepted by the seller contain additional information. From the seller's perspective, revealing this additional information is undesirable if it can be exploited by the dealers against her in the future, but withholding it may reduce the interest of the dealers in the trade. Our model explores this trade-off and shows how the seller's choice depends on various characteristics of the asset being traded and features of the market.

Our analysis highlights how the investor's choice of post-trade transparency affects market liquidity and the efficiency of bond allocations. We think that a natural setting where these results matter is electronic trading platforms. In such environments, it is possible to commit to a specific set of rules that would guide interactions between investors and dealers. As these platforms grow to take up an ever-increasing share of the secondary corporate bond market, the potential impact of their design increases as well. This paper stresses the importance of information embedded in the market participants' offers and points to post-trade disclosure rules as a potential tool to influence market performance.

References

Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael D. Grubb (2020) "Peaches, lemons, and cookies: Designing auction markets with dispersed information," *Games and Economic Behavior*, 124, 454–477.

Allen, Jason, Robert Clark, Brent Hickman, and Eric Richert (2023) "Resolving failed banks: Uncertainty, multiple bidding and auction design," *Review of Economic Studies*, Forthcoming.

Asquith, Paul, Thom Covert, and Parag Pathak (2019) "The effects of mandatory transparency in financial market design: Evidence from the corporate bond market," National Bureau of Economic Research.

Baldauf, Markus and Joshua Mollner (2023) "Competition and Information Leakage," *Journal of Political Economy*, Forthcoming.

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman (2006) "Market transparency, liquidity externalities, and institutional trading costs in corporate bonds," *Journal of Financial Economics*, 82 (2), 251–288.

Bloomberg Intelligence (2023) "Credit platforms shaking off signs of industry maturation," Available at https://www.bloomberg.com/professional/blog/credit-platforms-shaking-off-signs-of-industry-maturation/, [Last accessed May-23-2023].

Board, Simon (2009) "Revealing information in auctions: the allocation effect," *Economic Theory*, 38, 125–135.

Di Maggio, Marco, Amir Kermani, and Zhaogang Song (2017) "The value of trading relations in turbulent times," *Journal of Financial Economics*, 124 (2), 266–284.

Duffie, Darrell (2012) *Dark markets: Asset pricing and information transmission in over-the-counter markets*: Princeton University Press.

Edwards, Amy K, Lawrence E Harris, and Michael S Piwowar (2007) "Corporate bond market transaction costs and transparency," *The Journal of Finance*, 62 (3), 1421–1451.

Esó, Pèter and Balàzs Szentes (2007) "Optimal information disclosure in auctions and the handicap auction," *The Review of Economic Studies*, 74 (3), 705–731.

Fermanian, Jean-David, Olivier Guéant, and Jiang Pu (2016) "The behavior of dealers and clients on the European corporate bond market: the case of Multi-Dealer-to-Client platforms," *Market microstructure and liquidity*, 2 (03n04), 1750004.

Foucault, Thierry, Marco Pagano, and Ailsa Röell (2013) *Market liquidity: theory, evidence, and policy*: Oxford University Press.

Glode, Vincent, Christian C Opp, and Xingtan Zhang (2018) "Voluntary disclosure in bilateral transactions," *Journal of Economic Theory*, 175, 652–688.

Goldstein, Michael A, Edith S Hotchkiss, and Erik R Sirri (2007) "Transparency and liquidity: A controlled experiment on corporate bonds," *The Review of Financial Studies*, 20 (2), 235–273.

Hausch, Donald B. (1987) "An asymmetric common-value auction model," *The RAND Journal of Economics*, 611–621.

Hendershott, Terrence, Dmitry Livdan, and Norman Schürhoff (2021) "All-to-all liquidity in corporate bonds," Swiss Finance Institute Research Paper.

Hendershott, Terrence and Ananth Madhavan (2015) "Click or call? Auction versus search in the over-the-counter market," *The Journal of Finance*, 70 (1), 419–447.

Hendricks, Kenneth and Robert H. Porter (1988) "An empirical study of an auction with asymmetric information," *American Economic Review*, 865–883.

Kakhbod, Ali and Fei Song (2020) "Dynamic price discovery: Transparency vs. information design," *Games and Economic Behavior*, 122, 203–232.

Kong, Yunmi, Isabelle Perrigne, and Quang Vuong (2022) "Multidimensional auctions of contracts: An empirical analysis," *American Economic Review*, 112 (5), 1703–1736.

Kozora, Matthew, Bruce Mizrach, Matthew Peppe, Or Shachar, and Jonathan S Sokobin (2020) "Alternative trading systems in the corporate bond market," FRB of New York Staff Report 938.

Lewis, Ryan and Michael Schwert (2021) "The effects of transparency on OTC market-making," Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.

Madhavan, Ananth (1995) "Consolidation, fragmentation, and the disclosure of trading information," *The Review of Financial Studies*, 8 (3), 579–603.

Milgrom, Paul R. and Robert J. Weber (1982a) "A theory of auctions and competitive bidding," *Econometrica: Journal of the Econometric Society*, 1089–1022.

——— (1982b) "The value of information in a sealed-bid auction," *Journal of Mathematical Economics*, 10 (1), 105–114.

Myerson, Roger B (1981) "Optimal auction design," *Mathematics of Operations Research*, 6 (1), 58–73.

O'Hara, Maureen and Xing Alex Zhou (2021) "The electronic evolution of corporate bond dealers," *Journal of Financial Economics*, 140 (2), 368–390.

Pagano, Marco and Ailsa Röell (1996) "Transparency and liquidity: A comparison of auction and dealer markets with informed trading," *The Journal of Finance*, 51 (2), 579–611.

Pinter, Gabor, Chaojun Wang, and Junyuan Zou (2022) "Information chasing versus adverse selection," Bank of England Working Paper.

Reuters (2018) "MarketAxess allows Pimco to trade by its own rules," Available at https://www.reuters.com/article/us-pimco-marketaxess-idUSKBN1I528W, [Last accessed February-23-2023].

SIFMA (2016) "SIFMA Electronic Bond Trading Report: US Corporate & Municipal Securities," Report.

Vairo, Maren and Piotr Dworczak (2023) “What type of transparency in OTC markets?”, Available at SSRN 4431190.

Wang, Chaojun (2023) “The limits of multi-dealer platforms,” *Journal of Financial Economics*, 149 (3), 434–450.

Yueshen, Bart Zhou and Junyuan Zou (2022) “Less Is More,” Available at SSRN 4274063.

A Appendix

A.1 Proofs

Proof of Proposition 3. The results in Propositions 1 and 2 and Assumption 1 imply that the optimal bids of the buyer in the first period p_1^h and p_1^d are monotonically increasing with the private benefit B and are bounded by $\bar{v}$. Therefore, there is $\hat{B}$ for which $(1+\delta)\int_0^{p_1^h}(p_1^h - v)f(v)dv - (\bar{v} - \mathbb{E}\, v) = 0$. For any $B \geq \hat{B}$, we have $\Delta_s = (1+\delta)\int_0^{p_1^h}(p_1^h - v)f(v)dv - \int_0^{p_1^d}(p_1^d - v)f(v)dv \geq (1+\delta)\int_0^{p_1^h}(p_1^h - v)f(v)dv - (\bar{v} - \mathbb{E}\, v) \geq 0$ because a seller's per-period profit is bounded by $\bar{v} - \mathbb{E}\, v$. $\square$

Proof of Proposition 4. For distributions with thinner tails, the ratio $r(v) = \frac{f(v)}{F(v)}$, which is monotonically decreasing by Assumption 1, is steeper at the right tail. Because the buyer's optimal bids p_1^d and p_1^h are given by (5) and (16), for distributions with steeper $r(v)$, the distance between p_1^d and p_1^h is smaller. Thus, if this distance is sufficiently small, $|\Delta_1|$ is sufficiently close to zero while Δ_2 is bounded away from zero. As a result, $\Delta_s > 0$ for distributions with sufficiently thin tails. $\square$

Proof of Proposition 7. In this proof, we show that, if $v_2 = v_1$, the buyer's optimal bid $p_2^h(p_1)$ and profit $b_2^h(p_2^h(p_1))$ after the seller accepts the buyer's bid p_1 are the same as in the baseline model.

Because the results of Lemma 1 apply in the extended setup, if $v_2 = v_1$ and the bid p_1 is accepted, the buyer rationally infers that the seller's value v^a in the second period is distributed on $[\underline{v}, p_1]$. Thus, the buyer's profit $b_2^h(p_2)$ is given by the same function as in the baseline and the buyer's marginal benefit of increasing the bid p_2 is given by (11).

As in the baseline, if $p_1 \leq p_1^d$, by the monotonicity in the Assumption 1, we have $Bf(p_1) - F(p_1) \geq 0$. Therefore, the buyer's marginal benefit is positive at the upper boundary of the truncated distribution and the buyer optimally bids $p_2^h(p_1) = p_1$.

Consequently, if we conjecture that $p_1 \leq p_1^d$, the buyer's profit in the second period is

$$b_2^h(p_2^h(p_1)) = \int_0^{p_1} v \frac{f(v)}{F(p_1)} dv + (B - p_1). \tag{A.1}$$

Plugging this into the buyer's total profit evaluated at the first period, the buyer's marginal benefit of increasing the bid p_1 is given by (30). This is almost the same marginal benefit as the one for the first period trade in the case with disclosure (4) except there is now a additional marginal cost of $\delta\alpha F(p_1)$. Therefore, the buyer makes weakly more aggressive bids in the first period, i.e, $p_1^h \leq p_1^d$. Thus, the above conjecture is verified.

□

Proof of Proposition 8. The results of Proposition 2 can be extended to the case with $T > 2$ by backwards induction. In particular, the buyer's marginal benefit of increasing the bid p_t in period $t \in \{1, \dots, T\}$ is

$$(b_t^h)'(p_t) = Bf(p_t) - (1 + \delta + \cdots + \delta^{T-t})F(p_t) = Bf(p_t) - \left(\frac{1 - \delta^{T-t+1}}{1 - \delta}\right) F(p_t). \tag{A.2}$$

As in the baseline model, the marginal cost increases with the number of remaining periods $T - t$ while the marginal benefit does not depend on it.

Thus, the optimal bid of the buyer in the first period p_1^h is given by (33). Because the total marginal benefit (A.2) decreases with $T - t$, if p_1^h is accepted in any period $t \geq 1$, the buyer finds it optimal to submit the same bid at $t + 1$.

□

A.2 Multiple Bidders

Consider a setting where two bonds with the same underlying common value are sold in two consecutive auctions to several dealers. The dealers receive noisy signals about the value of the bond and compete in first-price sealed-bid auctions. If the auction holder chooses to conceal the bids, the participants, similar to our main model, have a benefit from losing in the first auction — they will have an informational advantage over the winner in the second auction. To the extent that this advantage may allow them to earn higher expected profit in the second auction than the winner, all the dealers in the first auction may reduce their quotes.

In this section, we would like to illustrate the idea that the seller may benefit from reducing information available to some of the bidders in the auction. In particular, we would like to show that, as a result of reduced information available to the bidders, it is possible that the seller earns higher profit while the difference between the profits of the bidders remains largely unaffected. It has been pointed out that, in settings with private values, the seller may prefer to keep information hidden from the bidders (see, e.g., Board, 2009). However, for our purposes, we would like to consider a pure common value auction. Analyzing models for common value auctions with asymmetric bidders is notoriously hard, so, instead of looking for the full set of conditions where the seller prefers not to inform the bidders, we focus on providing an illustration. To do this, we use a two-bidder model of Abraham et al. (2020).

Suppose two bidders, $i = 1, 2$, participate in an auction where the value of the underlying item v is common to both participants. Each bidder has a signal s_i that can take on one of the two values: either high or low. We normalize these values to either 0 or 1. Likewise, we normalize the possible values that v can take to 0 and 1. Signals are imperfectly informative of v, and one of the bidders has a more precise signal. Importantly, higher precision does not mean this bidder knows strictly more. Instead, we assume that signal s_1 is more strongly correlated with v, but s_2 still contains some information about v that is independent of s_1. We fully specify the joint distribution of the signals and the valuation in Table 1. Parameter $x > 0$ is such that $Cov(s_1, v) > 0$ does not depend on x while $Cov(s_2, v) > 0$ is decreasing in x. Furthermore, $Cov(s_1, s_2) > 0$ is also decreasing in x. Therefore, when x increases, the second bidder's signal gets less precise (so long as x is close enough to 0) and less correlated with the first bidder's signal. On the other hand, when $x = 0$, the bidders are symmetric, and their signals are positively correlated.

	$v = 0$		$v = 1$	
	$s_2 = 0$	$s_2 = 1$	$s_2 = 0$	$s_2 = 1$
$s_1 = 0$	$1/8 - x$	$1/8 + x$	0.02	0.11
$s_1 = 1$	$1/8$	$1/8$	0.11	0.26

Table 1: Joint distribution of bidder signals s_i and the underlying valuation v. In the table, x characterizes how less informed bidder 2 is compared to bidder 1.

Assuming the bidders receive signals and then submit bids in a first-price sealed-bid auction, there is a unique Nash equilibrium in monotone bidding strategies. We refer the reader to the original paper for reference. Here, this framework is used to highlight that when the seller can control the value of x, setting $x = 0$ is not necessarily a revenue-maximizing solution. Using the expression for revenue from Abraham et al. (2020), it is straightforward to find regions such that the seller prefers to set $x > 0$. We provide an illustration with seller profit in Figure 7 where we also include the difference between the bidders' expected profits. In this case, bidder 1 is better informed, so for $x > 0$ he is earning higher expected information rents. We can see from the figure that the benefits to the seller exceed the difference in the bidders rewards, suggesting that the incentives for the bidders to be relatively more informed are not necessarily as strong as the seller's incentives to conceal information.

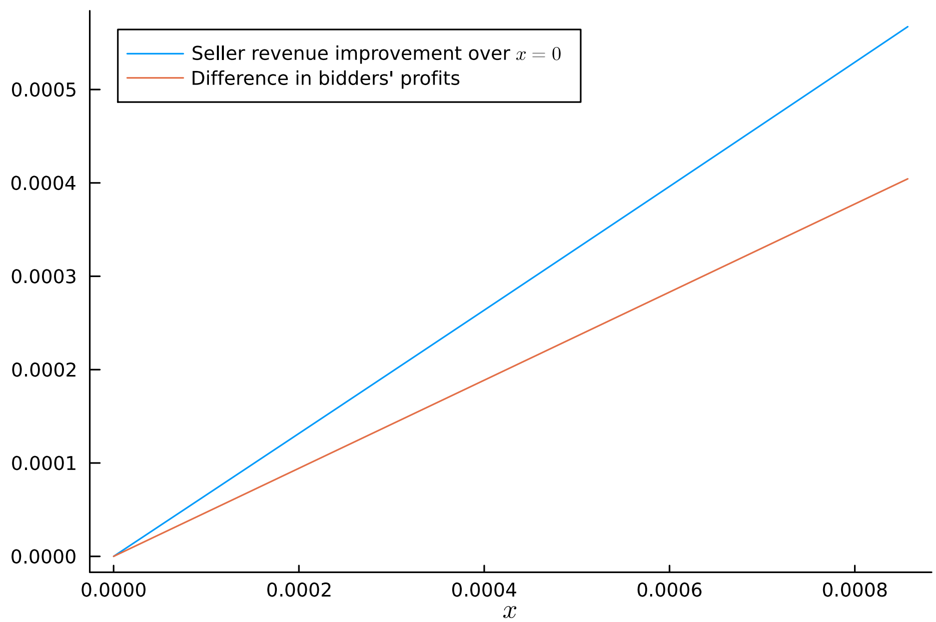

Figure 7: Seller revenue for an asymmetric common-value first-price auction. Based on the joint distribution of bidder signals and the underlying valuation in Table 1.

Is it always true that the seller prefers to introduce informational asymmetry between the bidders? The answer depends on the structure of the bidders' signals. Hausch (1987) discusses conditions under which the seller may wish to provide or withhold information from the bidders. The key idea is that bidder asymmetry could help to mitigate the winner's curse for the more informed bidder: she can feel more confident about the gap between her bid and the runner-up bid as being explained by the information asymmetry. Conversely, for symmetrically informed bidders, winner's curse can be stronger under some circumstances since it is definitely driven by having a more positive signal. Hausch (1987) provides some examples where the winner's curse is strong enough so that the seller chooses to withhold own, private information.

To further clarify the intuition in the auction setting, we describe a special case of a signal structure in which the seller makes the opposite choice. In particular, suppose the seller can decide whether the bidders have identical information or one of the bidders is strictly better informed than the other. If the bidders have identical information, they will have to compete away all the rents, similar to pure Bertrand competition. However, if one of the bidders is informed better, he can usually earn positive rents (Hendricks and Porter, 1988). Therefore, the seller would prefer not to hide information in this case. In the setup with repeated sales, this happens when the dealers do not receive any additional information about the bonds they are trading between the auctions. In this case, they can use each others' bids in the first auction to infer each others' signals, which would level the informational playing field and make them compete away the rents in the subsequent trades. The seller would then prefer not to hide the cover, and, in fact, would wish to reveal all bids submitted in the first auction.

A.3 Examples

A.3.1 Uniform Distribution

To illustrate the results of our main analysis, we have plotted graphs that use a numerical example in which v is uniformly distributed on $[0, 1]$. In this section, we present analytical results for this example.

Disclosed cover. For the uniform distribution on $[0, 1]$, $f(v) = 1$ and $F(v) = v$ on $[0, 1]$. Thus, in the case with a disclosed cover, the optimal bid of the buyer is implicitly given by

$$(b^d)'(p_1^d) = 0 \Leftrightarrow B - p_1^d = 0, \tag{A.3}$$

which yields the optimal bid:

$$p_1^d = \begin{cases} B & \text{if} \quad B \leq 1 \\ 1 & \text{if} \quad B > 1. \end{cases} \tag{A.4}$$

Thus, the seller's total profit is

$$s^d = \begin{cases} \frac{B^2}{2} & \text{if} \quad B \leq 1 \\ \frac{1}{2} & \text{if} \quad B > 1. \end{cases} \tag{A.5}$$

Hidden Cover. In the case with a hidden cover, the optimal bid of the buyer in the second period is implicitly given by

$$(b_2^h)'(p_2^h) = 0 \Leftrightarrow B - p_2^h = 0, \tag{A.6}$$

which yields the optimal bid:

$$p_2^h = \begin{cases} B & \text{if} \quad B \leq p_1^h \\ p_1^h & \text{if} \quad B > p_1^h. \end{cases} \tag{A.7}$$

The optimal bid of the buyer in the first period is implicitly given by

$$(b^h)'(p_1^h) = 0 \Leftrightarrow \frac{B}{1+\delta} - p_1^h = 0, \tag{A.8}$$

which yields the optimal bid:

$$p_1^h = \begin{cases} \frac{B}{1+\delta} & \text{if} \quad \frac{B}{1+\delta} \leq 1 \\ 1 & \text{if} \quad \frac{B}{1+\delta} > 1. \end{cases} \tag{A.9}$$

Thus, the seller's total profit is

$$s^h = \begin{cases} \frac{B^2}{2(1+\delta)} & \text{if} \quad \frac{B}{1+\delta} \leq 1 \\ (1+\delta)\frac{1}{2} & \text{if} \quad \frac{B}{1+\delta} > 1. \end{cases} \tag{A.10}$$

It can be seen that for any $B > \sqrt{1+\delta}$, the seller prefers to keep the cover hidden since

$s^h > s^d$ (see Figure 4).[18] Note that for any $B \in (\sqrt{1+\delta}, 1+\delta)$, the seller's profit in the first period is smaller in the case with a hidden cover, i.e., $\Delta_1 = s_1^h - s_1^d = \frac{B^2}{2(1+\delta)^2} - \frac{1}{2} < 0$, while the total seller's profit is larger i.e., $\Delta = s^h - s^d = \frac{B^2}{2(1+\delta)} - \frac{1}{2} > 0$ (see Figure 3).

Accounting for the seller's optimal choice of the regime with a disclosed cover or with a hidden cover, the buyer's total profit is

$$b^h = \begin{cases} \frac{B^2}{2} + \delta B & \text{if} \quad B < 1 \\ \frac{1}{2} + (B-1) + \delta B & \text{if} \quad B \in [1, \sqrt{1+\delta}] \\ \frac{B^2}{2(1+\delta)} + \delta B & \text{if} \quad B \in [\sqrt{1+\delta}, 1+\delta] \\ (1+\delta)(\frac{1}{2} + (\frac{B}{1+\delta} - 1)) + \delta B & \text{if} \quad B > 1+\delta. \end{cases} \tag{A.11}$$

Thus, the buyer's profit drops discontinuously at $B = \sqrt{1+\delta}$, with the difference being $\sqrt{1+\delta} - 1$ (see Figure 6).

A.3.2 Single-peaked Distribution

To illustrate the results of Proposition 4, we use a distribution with the density given by

$$f(v) = \begin{cases} \frac{1}{a} + v\left(4 - \frac{4}{a}\right) & \text{if} \quad v \in [0, \frac{1}{2}) \\ \left(4 - \frac{3}{a}\right) - v\left(4 - \frac{4}{a}\right) & \text{if} \quad v \in [\frac{1}{2}, 1]. \end{cases} \tag{A.12}$$

It is a symmetric distribution around $\frac{1}{2}$ with $f(0) = f(1) = \frac{1}{a}$ and $f(\frac{1}{2}) = 2 - \frac{1}{a}$. The parameter a controls the steepness of $r(v)$ at the right tail, i.e., higher a corresponds to thinner tails. The uniform distribution is a special case for $a = 1$.

In the figures below, we plot the buyer's optimal bids and the seller's profits for our two main cases, when the cover is disclosed and when it is hidden, as well the difference in the seller's profits between the two cases.

[18] In the case of the uniform distribution, the hidden cover is more profitable for the seller only if the buyer makes the least aggressive bids in the case with disclosure, i.e., only if $p_1^d = \bar{v} = 1$. However, this is not a necessary condition as it can be seen from the example of the next section.

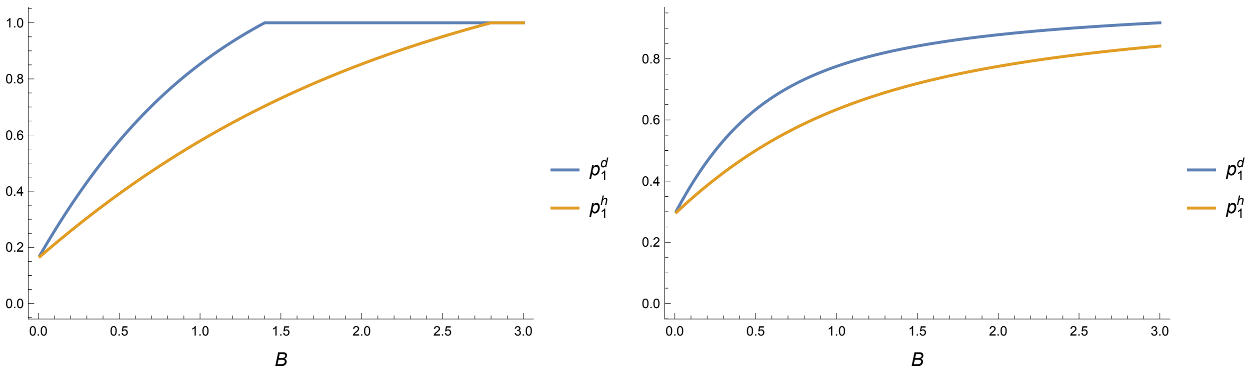

Figure 8: The buyer's bids in the first period in the cases with a disclosed cover, p_1^d, and a hidden cover, p_1^h. In the graph, we consider a setting with $f(v) = \left[\frac{1}{a} + v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(0 \leq v \leq \frac{1}{2}) + \left[\left(4 - \frac{3}{a}\right) - v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(\frac{1}{2} < v \leq 1)$, $B = 1$, $\delta = 1$ for $a = 1.4$ (left) and $a = 100,000$ (right).

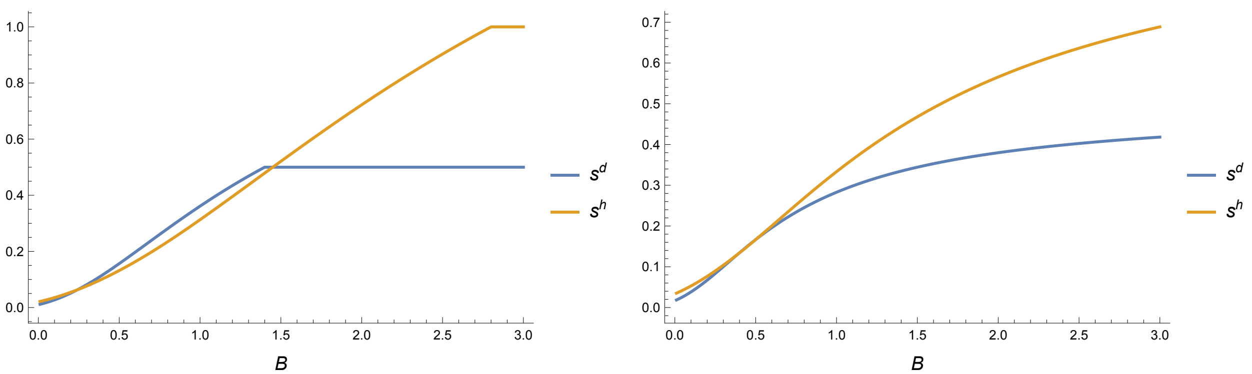

Figure 9: The seller's total profit in the cases with a disclosed cover, s^d, and a hidden cover, s^h. In the graph, we consider a setting with $f(v) = \left[\frac{1}{a} + v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(0 \leq v \leq \frac{1}{2}) + \left[\left(4 - \frac{3}{a}\right) - v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(\frac{1}{2} < v \leq 1)$, $B = 1$, $\delta = 1$ for $a = 1.4$ (left) and $a = 100,000$ (right).

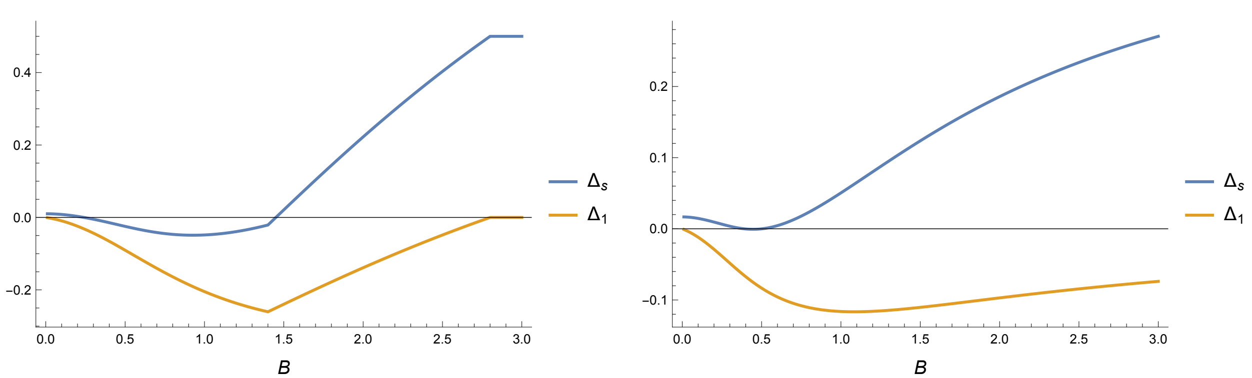

Figure 10: The difference in the seller's total profit, Δ_s and the seller's first-period profit, Δ_1, between the cases with a hidden cover and a disclosed cover. In the graph, we consider a setting with $f(v) = \left[\frac{1}{a} + v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(0 \leq v \leq \frac{1}{2}) + \left[\left(4 - \frac{3}{a}\right) - v\left(4 - \frac{4}{a}\right)\right] \mathbb{1}(\frac{1}{2} < v \leq 1)$, $B = 1$, $\delta = 1$ for $a = 1.4$ (left) and $a = 100,000$ (right).