Optimal Transparency in Over-the-Counter Markets*

Maren Vairo† and Piotr Dworczak‡

August 14, 2026

Abstract

Financial over-the-counter markets often lack both pre-trade transparency (visibility of current quotes) and post-trade transparency (disclosure of past trades), which raises concerns about their efficiency. We propose a dynamic trading model with adverse selection to study both forms of transparency. We find that post-trade transparency improves welfare relative to the opaque market but is dominated by pre-trade transparency; moreover, in markets that are already pre-trade transparent, introducing post-trade transparency does not generate further gains and can sometimes reduce welfare. To test the robustness of these conclusions, we use a mechanism-design approach to characterize the *optimal* trading mechanism in our framework. The optimal outcome can be implemented through a trading protocol that offers pre-trade but not post-trade transparency.

Keywords: Over-the-counter markets, pre-trade transparency, post-trade transparency, adverse selection, efficiency

*We thank the editor and three anonymous referees, as well as Philip Bond, Belén Jerez, Benjamin Lester, Vincent Maurin, Josh Mollner, Humberto Moreira, Alessandro Pavan, Harry Pei, Michael Peters, the Trading Policy team at the UK Financial Conduct Authority, and several seminar audiences for helpful comments and suggestions. Dworczak gratefully acknowledges the support of the Alfred P. Sloan Fellowship. Eric Chen and Pedro Aldighieri provided excellent research assistance. The paper was previously circulated under the title "What Type of Transparency in OTC Markets?".

†Wharton School, University of Pennsylvania.

‡Department of Economics, Northwestern University.

1 Introduction

Many financial assets—including corporate bonds, swaps, currencies, non-standard derivatives, real estate, and consumer products—are primarily traded in over-the-counter (OTC) markets. *Opacity* has been traditionally one of the key features setting OTC markets apart from traditional centralized exchanges (Duffie, 2012). Traders in financial over-the-counter markets must search for counterparties before they transact (Duffie et al., 2005). Often, they must request a quote from an OTC dealer and decide whether to accept it before they can learn what prices other dealers may be offering (Zhu, 2012). This unawareness of the currently available quotes can be described as lack of *pre-trade transparency*. Similarly, in most OTC markets, traders do not observe each other's transactions, and hence cannot rely on data about past prices to inform their trading strategies or refine their estimates of asset values. This feature is generally referred to as lack of *post-trade transparency*.

The opacity of financial OTC markets has long raised concerns among policymakers, which were exacerbated by the 2007-2008 financial crisis. In the US, a number of regulations have increased both the pre- and post-trade transparency of certain OTC markets. Already in early 2000s, the SEC instituted the Trade Reporting and Compliance Engine (TRACE) in the US corporate bond market.[1] Under TRACE, broker-dealers who are FINRA members are required to disseminate price and volume information about their trades, effectively in real time. This results in considerably increased post-trade transparency in markets covered by TRACE. By providing access to data about past prices, the regulator intended to promote market efficiency by "leveling the playing field." More recently, in the aftermath of the so-called Libor scandal, regulators around the world took actions to support robust financial benchmarks (e.g., the replacement of LIBOR with SOFR in the US) that are also an important source of post-trade transparency through public dissemination of *average* prices.[2] The 2010 Dodd-Frank Act, on the other hand, promoted pre-trade transparency in swap markets by introducing Swap Execution Facilities (SEFs). In contrast to traditional search markets, SEFs were designed to allow customers to request quotes from multiple dealers at the same time. Transparency regulations have also significantly increased in the European Union. The revised Markets in Financial Instruments Directive (MiFID II), implemented in 2018, promotes both forms of transparency by requiring a minimum volume of transactions to be conducted electronically through regulated platforms and mandating trade reporting for most transactions. Further modifications to the MiFID II regulatory framework are currently under discussion (FCA, 2024).

[1] See, for example, Asquith et al. (2019).

[2] See, among others, Hou and Skeie (2013), Duffie and Stein (2015), and Duffie et al. (2017).

This paper addresses the problem of *optimal* design of transparency in a financial over-the-counter market. We propose a simple dynamic model of trading under adverse selection in which asymmetrically informed dealers compete to trade with uninformed customers. We investigate how different forms of pre- and post-trade transparency affect the equilibrium level of welfare. In the first part of the paper, we compare the performance of the market under several regimes created by the regulation discussed above. In the second part, we use mechanism design to identify the optimal market design in a broad class of dynamic trading mechanisms. We view the main contribution of the paper as providing the first theoretical framework that accommodates *both* pre- and post-trade transparency, resulting in novel insights that could guide future regulation of OTC markets.

Our model features two dealers, who are long-run players competing for the order flow, and an infinite sequence of customers who arrive to the market desiring to buy or sell one unit of the asset. The asset has a common-value component; a customer's value differs from the common value by a constant private-value component reflecting, for example, a liquidity need. Thus, there is a constant, strictly positive gain from trade between customers and dealers, and efficiency requires trade to happen in every period. An adverse selection problem precludes this efficient outcome: Dealers observe private signals about the common value of the asset, while customers are initially uninformed.

We ask how different transparency regulations affect the equilibrium probability of trade (which is equivalent to a standard measure of allocative efficiency in our framework). We model the traditionally opaque OTC market by assuming that the customer does not have access to any information about past trades, and can only find out about dealers' quotes by visiting them. Motivated by the popular request-for-quote (RFQ) protocol featured on many trading platforms, we model a fully pre-trade transparent market as a first-price auction between the dealers.[3] We model post-trade transparency by letting customers who arrive to the market observe the (anonymized) history of transactions. Finally, a fully transparent market has both pre- and post-trade transparency.

Our framework endogenizes the influence that these different transparency regimes have on equilibrium learning and information asymmetry. Customers are initially uninformed but can learn about the asset value from dealers' quotes in three ways, depending on the setting: *(i)* by requesting quotes (in the opaque regime), *(ii)* by observing dealers' quotes in the auction (in pre-trade transparent regimes), and *(iii)* by inspecting dealers' past executed quotes (in post-trade transparent regimes). We show that these three information channels differ significantly in terms of their equilibrium consequences for dealers' pricing strategies, and

[3] We assume that dealers observe each other's quotes after the auction, which turns out to be important for our results—see Section 3.3 for a detailed discussion.

hence welfare. In opaque markets, dealers' informational advantage creates an adverse selection problem for the customer, and supporting an equilibrium requires quotes that are often rejected.[4] In pre-trade transparent markets, in addition to the customer's adverse-selection problem already present under opacity, direct competition exposes uninformed dealers to adverse selection as well. To avoid losses when competing against informed dealers, uninformed dealers submit more conservative quotes, which reduces efficiency. Finally, in post-trade transparent markets, since customers observe executed transaction prices, dealers have an incentive to manipulate future customers' beliefs by quoting as if they had more favorable information about the asset value. Discouraging this behavior in equilibrium requires that such quotes be accepted less often, which again reduces the equilibrium probability of trade.

Unsurprisingly—given the presence of signaling motives and adverse selection—our model features equilibrium multiplicity. We impose only a few intuitive equilibrium refinements motivated by practical considerations, and compare different regimes using two robust dominance notions. Specifically, one market regime is said to *weakly welfare-dominate* another one if it has higher expected welfare in both the best and the worst equilibrium; *strong welfare dominance* requires that the worst equilibrium of one regime has higher expected welfare than the best equilibrium of another regime. Welfare comparisons depend on the value of a parameter that governs the severity of adverse selection. This parameter is higher in markets where dealers are more likely to possess private information relevant to the asset's value.

Figure 1.1 summarizes our main findings in the first part of the paper. Post-trade transparency improves upon the opaque market; however, the improvement is only in the weak sense. Pre-trade transparency welfare-dominates post-trade transparency, and the improvement is in the strong sense in markets with a high degree of adverse selection (the case depicted in the figure). Adding post-trade transparency to a pre-trade transparent market can only reduce welfare: The worst equilibrium stays the same, while the best equilibrium may become less efficient.

A high-level intuition for these results is as follows. Post-trade transparent markets become efficient once past quotes reveal the value of the asset. However—by definition—post-trade transparency cannot positively affect efficiency immediately after new information arrives to the dealers, since at least one informative transaction must take place before information gets disseminated to other market participants. Moreover, disclosure of past trades implies that dealers can attempt to manipulate market beliefs by trading at distorted prices, which further constrains efficiency prior to disclosure of a transaction.

[4] An efficient pooling equilibrium is ruled out in our analysis by assuming that private gains from trade are not too large compared to the uncertainty about the common value component.

Figure 1.1: Welfare ranking across trading regimes under severe adverse selection

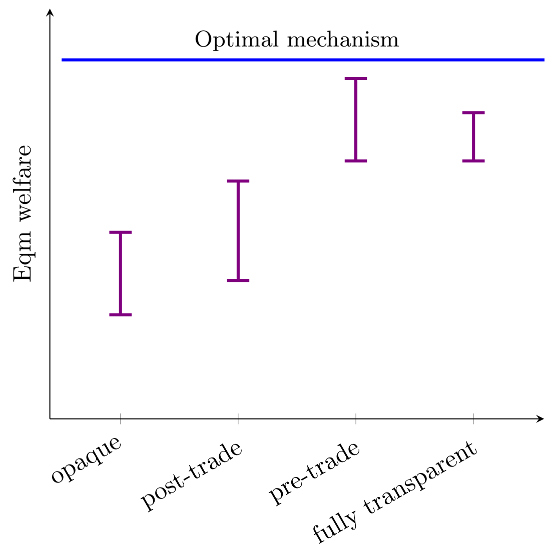

Pre-trade transparency is overall more efficient than post-trade transparency for two reasons: First, direct competition between dealers makes them offer attractive quotes, which are accepted by customers with high probability; second, dealers' information becomes symmetric after they observe each other's quotes, making subsequent trades fully efficient. Both design aspects—simultaneity of quotes and the ability of dealers to see each other's quotes—are crucial for pre-trade transparency to deliver its benefits. In particular, we argue that welfare would be low in a sequential-search market even if it were effectively free for customers to request quotes from dealers. This highlights the point that the details of the protocol matter for the effects of competition in the face of adverse selection.

The detrimental effect of adding post-trade transparency to a pre-trade transparent regime follows from the observation that once dealers are symmetrically informed and compete directly, customers do not *need* to be informed to make efficient trading decisions. An attempt to make customers more informed by disclosing past transactions "backfires" as it creates incentives for dealers to manipulate the beliefs of future customers. Preventing manipulation requires capping the probability of executing quotes that induce market beliefs favorable to the manipulator, which reduces welfare.

Although our equilibrium analysis covers several transparency regimes, it does not encompass the full range of trading mechanisms used in practice. Regulation leaves scope for market participants to experiment with the designs of trading mechanisms. For example, in the US swap market, traders can choose from several SEFs that differ in their pre- and post-trade transparency, as well as other features (SIFMA, 2016). In the second part of the paper, we turn to the problem of *optimal* market design. We characterize the welfare-maximizing dynamic mechanism subject to natural constraints, such as market clearing, budget balance, incentive compatibility and individual rationality.[5] We also require the mechanism to be stationary: Trading rules may depend on the relevant market conditions but not on calendar time or payoff-irrelevant transaction history.[6]

We find that the first-best outcome (trade with probability one) can be implemented for a wide range of parameter values and that the optimal mechanism outperforms the transparency regimes discussed in the first part of the paper.[7] To gain intuition for the properties of the optimal mechanism, we construct an indirect trading mechanism that implements the optimal outcome in the best-case equilibrium. The trading protocol is static (it only conditions on current-period quotes), and resembles a second-price auction combined with an ex-post bargaining stage. The mechanism features no post-trade transparency, that is, it discloses no information to traders who did not participate in the auction. Intuitively, in a second-price auction, conditional on winning, a less-informed dealer trades at the more-informed dealer's quote, which alleviates the adverse selection problem. The post-auction bargaining stage corrects a potential inefficiency associated with the fact that the customer—say, a buyer—does not trade at the lowest available quote in a second-price auction. It is possible that the buyer's updated estimate of the value of the asset (conditional on observed quotes) lies strictly below the second-lowest quote but above the quote of the winning dealer. In this case, the post-auction bargaining stage allows the winning dealer and the customer to negotiate a price that makes trade profitable for both. This last finding explains why it may be efficient for platforms to allow for bilateral negotiations on top of offering an auction-like protocol; many, though not all, SEFs offer such a feature.[8]

The rest of the paper is organized as follows. We review the related literature next. In Section 2, we introduce the model. In Section 3, we present our main results on welfare properties of equilibria under different transparency regimes. In Section 4, we study the optimal mechanism. In Section 5, we discuss extensions and limitations of the model, focusing on how they shape policy implications of our results. All proofs are collected in the appendix.

[5] We impose an ex-post version of incentive constraints to circumvent the well-known issue with mechanism design models with correlated information, as explained in Crémer and McLean (1985).

[6] Fully general mechanisms could use past history to construct elaborate punishment schemes to discipline the dealers, but we view such mechanisms as impractical.

[7] This is not obvious a priori because ex-post implementation is a stronger solution concept than the equilibrium notion used in the first half of the paper.

[8] MarketAxess, DelphX, TradeWeb, and Bloomberg platforms serve as an example; see SIFMA (2016).

1.1 Related literature

There is a sizable theoretical literature devoted to the analysis of transparency of financial OTC markets. Our primary contribution is to propose a framework that allows for a tractable analysis of both pre- and post-trade transparency (as well as varying other aspects of the market design). This leads to novel insights—for example, about the interaction of the two types of transparency—that would not be possible to obtain otherwise. Additionally, by studying the design of the optimal trading mechanism within this framework, we join a nascent literature applying mechanism design to understand optimal trading protocols in financial markets.[9]

To the best of our knowledge, the theoretical literature has focused on the effects of introducing a single type of transparency at a time. For example, Asriyan et al. (2017), Duffie et al. (2017), Back et al. (2020), and Ollar et al. (2021) analyze the consequences of adding various forms of post-trade transparency to an opaque market, while Riggs et al. (2020) and Glebkin et al. (2023) study the effects of pre-trade transparency (introduction of auction-like mechanisms).[10] Each of these papers adopts a different model of trading, so comparing their results and understanding the interaction of the forces they identify is challenging. Our paper offers a simple framework in which the effects of pre- and post-trade transparency can be analyzed jointly. Throughout the analysis, we relate our conclusions to the results in these papers and explain which new insights emerge due to the unified approach to transparency.

Our model is most similar to the framework introduced by Kakhbod and Song (2020) (see also Kakhbod and Song, 2022). They find that post-trade transparency can have an adverse effect on price informativeness: In a repeated interaction between traders, disclosure of past prices allows for richer punishments for off-path behavior, which in turn makes it possible to support opaque pooling equilibria for a larger set of parameters. Thus, paradoxically, more post-trade transparency can lead to less information being disclosed on equilibrium path.[11] Our paper is complementary as it studies a different objective (allocative rather than price efficiency), focuses on separating equilibria (we rule out pooling equilibria with an assumption about model parameters), and analyzes pre-trade transparency (on top of post-trade transparency). In a different framework (one where it is the uninformed party

[9] Antill and Duffie (2020) and Chen (2024) use mechanism design to find the optimal trading protocol under the assumption that traders observe independent signals; Baldauf and Mollner (2020) apply mechanism-design tools to understand optimal design of exchanges in the presence of high-frequency trading; Loertscher et al. (2022) solve a dynamic mechanism-design problem to characterize optimal market thickness.

[10] Relatedly, Rostek and Wu (2021) model pre-trade access to price information in fragmented markets by allowing traders' demands in one venue to condition on simultaneously determined prices in other venues.

[11] Kakhbod and Song (2020) show that this increased opacity may be beneficial for market liquidity. A similar force is uncovered by Bergemann and Hörner (2018) in the context of a repeated first-price auction.

who makes offers), Kaya and Roy (2024) show that the effects of post-trade transparency on equilibrium learning about the asset value crucially depend on the competition in the market.

More broadly, our framework connects to many classical literatures on dynamic adverse selection, in particular in the context of search (e.g., Wolinsky, 1990, Blouin and Serrano, 2001, Lauermann and Wolinsky, 2016, Lester et al., 2019) and bargaining (e.g., Hörner and Vieille, 2009, Fuchs et al., 2016, Kim, 2017, Kaya and Kim, 2018). These literatures studied the influence of information disclosure on allocative efficiency and information aggregation through prices. While certain forces that drive our welfare results are familiar from these models, to the best of our knowledge, pre- and post-trade transparency have not been studied together within a single framework. In Appendix OA.1, we review these literatures in more detail and explain both the similarities and differences to our setting.

On the empirical side, a series of papers have studied the effects of increased post-trade transparency, through the introduction of TRACE, in the US corporate bond market (Edwards et al., 2007, Bessembinder and Maxwell, 2008, Asquith et al., 2019). The findings in these papers indicate that post-trade transparency led to a reduction in transaction costs. However, Asquith et al. (2019) find that this reduction did not necessarily translate into higher trading activity. They also document that the effect of TRACE on market liquidity differs considerably, both in sign and in magnitude, across bond types. Aquilina et al. (2022) find that increasing the informativeness and precision of benchmarks in the swaps market led to improvements in the quality of dealer-trader matches, thereby increasing the volume of welfare-enhancing trades. For the European market, Meli et al. (2024) find ambiguous effects of MiFID II on market liquidity. Overall, these empirical findings are consistent with the results in our paper predicting that the effects of post-trade transparency tend to be positive (relative to opacity) but may depend on the details of the market (and possibly on equilibrium selection).

With regard to pre-trade transparency, Hendershott and Madhavan (2015) analyze the effects of electronic trading, through periodic first-price auctions, on the corporate bond market. They document that the resulting increase in competition between dealers led to more attractive prices and higher market liquidity, which is in line with our findings. Riggs et al. (2020) study pre-trade transparency (using SEF trading data) but focus on the optimal behavior of customers; they show that customers contact only a few dealers and that some dealers refrain from responding to requests. This last observation is consistent with our finding that uninformed dealers may fail to trade in the pre-trade transparent market.

2 Model

2.1 The trading environment

We consider a trading environment with discrete time and infinite horizon, $t = 0, 1, \ldots$, a single risky asset, two long-lived dealers, and a sequence of short-lived customers.

Asset. The risky asset has a common value $v_t \in \{-1, 1\}$ at time t. At time $t = 0$, v_0 is drawn from the uniform distribution—i.e., $\Pr(v_0 = -1) = \Pr(v_0 = 1) = 1/2$. In every subsequent period t, the value of the asset is either "persistent" ($v_t = v_{t-1}$) or the value of the asset "resets," that is, it is redrawn from the uniform distribution. We interpret the reset of the value of the asset as arrival of exogenous payoff-relevant news to the market.

Persistence in the asset value is necessary to analyze the effects of post-trade transparency but it complicates the analysis in an infinite-horizon model with learning: We model the evolution of the value of the asset in a stylized and parsimonious way by assuming that the asset value can be persistent over at most two periods. (In Section 5, we comment on why our insights would continue to hold qualitatively without that assumption.) Specifically, there is an underlying Markov chain $Z_t \in \{0, 1\}$, with $Z_t = 0$ representing the event that v_t gets redrawn in period t (news arrives to the market), and $Z_t = 1$ representing the event that $v_t = v_{t-1}$ (no news arrives to the market). Conditional on news arriving to the market at time $t - 1$, there is probability $\gamma \in (0, 1)$ that no news will arrive in period t. Formally, $\Pr(Z_t = 1 | Z_{t-1} = 0) = \gamma$. However, news arrives for sure after a period with no news—i.e., $\Pr(Z_t = 1 | Z_{t-1} = 1) = 0$. Thus, γ is a measure of persistence of the asset's value. Under stationarity assumptions that we will impose, equilibrium behavior needs to be characterized only conditional on one of the two possible states—the "reset state" $Z_t = 0$ and the "persistent state" $Z_t = 1$—which greatly simplifies the analysis.[12]

Dealers. Two dealers, labeled $i \in \{A, B\}$, compete in the market for the asset by providing quotes to customers, according to a trading protocol that will be specified later. Both dealers value each unit of the asset at v_t at time t, and face no inventory or short-selling constraints. They discount the future through a common discount factor $\delta \in (0, 1]$. Formally, if in each period t a dealer sells $x_t^s \in \mathbf{R}_+$ units of the asset at a price a_t and buys $x_t^b \in \mathbf{R}_+$ units at a price b_t, her intertemporal payoff is $\sum_{t=0}^{\infty} \delta^t (x_t^s(a_t - v_t) + x_t^b(-b_t + v_t))$.

Dealers differ in the information they have about the asset value. Conditional on news arriving to the market in period t ($Z_t = 0$), each dealer privately observes a signal realization at the beginning of that period. Conditional on v_t, the two signals are independent, and fully

[12] While we find the infinite-horizon version of our model more natural to interpret, we note that all results would continue to hold in a simpler two-period model with a fully persistent state but with an exogenous probability $1 - \gamma$ that the game ends (e.g., the asset is liquidated) after the first period (see Appendix A.1).

reveal the value v_t with conditional probability $q \in (0,\, 1)$; with the remaining probability $1-q$, the signal conveys no information about v_t. If no news arrive to the market in period t ($Z_t = 1$), dealers do not observe any new signal realizations. We refer to a dealer who observed the uninformative signal realization as the "uninformed type;" we refer to a dealer who observed the revealing signal realization as the "informed type;" we also use the "high type" (resp., "low type") to refer to an informed dealer who knows that $v_t = 1$ (resp., $v_t = -1$). We index these three dealer types by $\sigma \in \{l, u, h\}$.

Our assumptions about the information structure imply that while news arrival to the market (state Z_t) is common knowledge, the dealers may not succeed in "deciphering" the implications of the news for the asset value. Alternatively, as in the models of Budish et al. (2015) and Baldauf and Mollner (2020), dealers may experience random latency between the arrival of news to the market and the time this information is reflected in their quotes. A key feature of the information structure is that a dealer may be less informed than her competitor in some periods.[13]

Customers. In each period t, a short-lived customer arrives to the market, desiring to trade one unit of the asset. The customer is a buyer or a seller with equal probability. The value for the period-t customer is $v_t + \Delta$ if she is a buyer and $v_t - \Delta$ if she is a seller, where $\Delta \in (0,\, 1)$ is a known constant, reflecting an immediacy need or a similar customer-specific benefit for trading the asset. Therefore, the payoff of a period-t customer-buyer who trades $x \in \{0, 1\}$ units at a price a is $x(-a + v_t + \Delta)$, while the payoff of a customer-seller trading $x \in \{0, 1\}$ units at a price b is $x(b - v_t + \Delta)$. By symmetry across these two scenarios, we can without loss of generality focus on the case when the current-period customer wants to buy the asset in all of our discussions.[14]

Customers possess no exogenous private information. (In Appendix OA.2, we extend our model to allow customers to have private information about the gains from trade; our conclusions remain unchanged as long as the amount of private information is small.) We assume that customers know the current state Z_t when they arrive to the market; that is, customers observe the event of news arrival but lack the expertise to interpret it. In light of this assumption, we can think of the parameter q as measuring the informational advantage of dealers over customers.

[13] Observe that if both dealers are informed, they must agree about the value of the asset (we cannot have a low type and a high type in the same period). This is a consequence of assuming a "truth-or-nothing" signal, which is commonly used in the literature on verifiable information following Dye (1985). In the context of financial markets it was used, for example, by Baldauf and Mollner (2020).

[14] Given that the asset value is distributed symmetrically around zero, our analysis of the customer-buyer case can be translated to the customer-seller case by switching the roles of high and low dealer types and changing the sign of prices. Note, however, that dealers face uncertainty about the desired direction of trade of future customers.

2.2 Market protocols

In all the protocols that we study, in every period, each dealer provides a single quote to the customer upon request. Upon receiving a quote, the customer can accept or reject it. The differences between trading regimes lie in who can observe a given quote and when.

1. In the **opaque market**, the customer can only see the quote of one of the dealers. That quote is not observable to anybody else. Requesting a quote from the second dealer is assumed to be prohibitively costly.

2. In the **post-trade transparent market**, the protocol is the same as in the opaque market except that executed quotes are publicly revealed: When a dealer trades with a customer in period t, the trade direction and price are observable to all traders from period $t + 1$ onward. Identities of trading parties are not disclosed.

3. In the **pre-trade transparent market**, the customer can see both quotes simultaneously. The quotes are also observable to the quoting dealers (but they are not observed by customers in consecutive periods.).

4. In the **fully transparent market**, the protocol is the same as in the pre-trade transparent market except that executed quotes are publicly revealed: When a dealer trades with a customer in period t, trade direction and price are observable to all traders from period $t + 1$ onward.

Our opaque market captures a trading environment in which customers receive no public information about past or current prices. Moreover, search costs are large enough that the customer never visits more than one dealer.

Our modeling of the post-trade transparent market is inspired by regulation such as the introduction of TRACE in the US corporate bond market (see, for example, Asquith et al., 2019). The main consequence of such transparency is that traders can condition their pricing and trading strategies on past (anonymized) transaction data. Such regulation is often motivated as an attempt to "level the playing field" for smaller investors by giving them access to information that was only available to dealers in opaque markets.

Our modeling of pre-trade transparency is inspired by regulation centralizing opaque OTC markets by requiring all trades to be executed on trading platforms. A prominent example is the introduction of Swap Execution Facilities that must be used in the US swap market following the Dodd-Frank Act. The goal of the regulation is to increase competition between dealers by providing the requester with *simultaneous* access to all quotes.

The fully transparent market has almost all features of a centralized trading protocol, typical for traditional stock exchanges. Traders have simultaneous access to all available quotes and observe past transaction data. However, even in this case, our framework retains one key feature of an OTC market: the privileged position of dealers in the trading network.

In the earlier working paper Vairo and Dworczak (2024), we also studied an opaque market with *low search costs*, where the customer could "shop around" for the best quote by visiting a second dealer at an arbitrarily small cost after seeing the first dealer's quote. We showed that sequential search exacerbates the adverse selection problem, so the low-search-cost market is typically *less* efficient than the high-search-cost market.[15] The possibility of adversarial interplay between search and adverse selection has been recognized in the literature (e.g., Wolinsky, 1990; Blouin and Serrano, 2001; Lauermann and Wolinsky, 2016; Lester et al., 2019). Since our main welfare results are robust to the low-search-cost case, we focus on the high-search-cost regime in this paper.

2.3 Solution concept

An equilibrium is described by players' strategies and their beliefs at every information set. Specifically, under each of the protocols that we study, the strategy of dealer $i \in \{A, B\}$ consists of a history-contingent distribution over quotes, and the period-t customer's strategy of her decision to accept one of the available quotes or exit the market without trading. Our solution concept is symmetric stationary perfect Bayesian equilibrium with separating strategies for the dealers. By stationary equilibrium, we mean that—conditional on the hierarchy of beliefs about dealers' types—dealers' and customers' strategies can depend on calendar time t only through the (publicly observed) state Z_t.[16] By a separating strategy for a dealer we mean that, on equilibrium path, the quote of dealer i at time t should uniquely pin down dealer i's belief about the value of the asset at time t. By symmetric equilibrium, we mean that the two dealers must follow the same strategy if their belief hierarchy is symmetric, and that a customer must treat the two dealers symmetrically if they enter her belief hierarchy symmetrically.

The restriction to equilibria with separating strategies for dealers is our strongest assumption. While $\Delta < 1$ rules out efficient pooling equilibria (regardless of the regime), there may exist semi-pooling equilibria. Aside from tractability considerations, which we discuss in Section 3, we choose not to consider semi-pooling equilibria for two reasons. First, prices

[15] Customers choose to search only after observing a quote that gives them an informational advantage over the uninformed dealer; this exposes the uninformed dealer to strong adverse selection and often results in no trade in equilibrium. The mechanism is reminiscent of the adverse selection problem faced by market makers in classic models such as Kyle (1985) and Glosten and Milgrom (1985), with the difference that adverse selection arises endogenously from customers' ability to gather information through search when search costs are low. The mechanism is also consistent with empirical evidence from the corporate bond market offered by Kargar et al. (2024) who document that search may be quite inefficient.

[16] A similar refinement, under the name of *weak Markov property*, is used by Schmidt (1993).

in financial markets play a dual role: They guide trading decisions in the market but they also serve as signals for actors in the real economy (Fama and Miller, 1972). By focusing on separating equilibria, we keep the second (unmodeled) welfare channel invariant across all regimes. If some semi-pooling equilibria increased welfare in our model relative to separating equilibria, we would be implicitly introducing a trade-off since these equilibria would provide less information to the (unmodeled) outsiders relying on financial-market prices to make real-economy decisions. Second, the focus on separating equilibria seems natural in the context of transparency regulation, since the goal of the regulation, by definition, is to increase the amount of information revealed to the market; if market participants respond to the regulation by switching to an uninformative pricing strategy, the regulation does not achieve its goal (a point made elegantly by Kakhbod and Song, 2020).

Studying stationary equilibria—beyond providing a significant analytical convenience—allows us to focus on how transparency affects trading through its impact on beliefs about the asset value. It is well known that transparency can also affect dynamic trading via forces familiar from the study of repeated games—conditioning continuation play on payoff-irrelevant signals expands the set of equilibrium outcomes in the current period.

Our environment induces a signaling game. In the absence of equilibrium refinements on off-path beliefs, economically unnatural outcomes can be sustained by equilibria belonging to the class defined above. Equilibrium refinements that uniquely select a separating equilibrium have been widely used in the literature on signaling (Banks and Sobel, 1987, Cho and Kreps, 1987) and on adverse selection (Guerrieri et al., 2010, Guerrieri and Shimer, 2014). However, these refinements are not readily applicable in our environment which is dynamic and features multiple informed players. In Appendix OA.3, we apply the D1 criterion of Banks and Sobel (1987) to the opaque regime (which is effectively static) and show that it uniquely selects the most efficient separating equilibrium, thereby suggesting that semi-pooling equilibria exhibit undesirable properties in our framework—consistent with typical findings in signaling games.

Instead of attempting to select a unique equilibrium, we impose relatively mild refinements on customers' off-path beliefs that discipline our analysis and rule out certain outcomes that we view as economically unintuitive. First, we assume a version of the *never-dissuaded-once-convinced refinement* commonly used in bargaining models with incomplete information (see Osborne and Rubinstein, 1990). Specifically, we restrict the support of customers' beliefs after observing an off-path event to be a subset of the support of her beliefs before observing that event. Second, we assume a *monotonicity refinement*: For an off-path quote b, *(i)* the customer's expectation of the asset value must be monotone in b; and *(ii)* whenever possible, the customer must attribute off-path quotes to dealer types that would not make a loss in case of trade.[17] Finally, we impose an additional *unprejudiced-belief refinement* (which only has bite in pre-trade transparent regimes): Whenever a pair of simultaneous quotes is such that exactly one of them is consistent with on-path play, and interpreting the pair through the on-path quote is compatible with the first two refinements, beliefs must be consistent with the quote that is on path.[18]

We offer an extended discussion of these refinements (including their formal definitions) in Appendix OA.3. Here, we make a few comments about their meaning and implications in the context of our trading game. First, the never-dissuaded-once-convinced refinement allows us to capture the intuitive property that trading should be efficient once all traders have the same information about the asset value. Without the refinement, even full disclosure of the (persistent) asset value in period t may not prevent trade breakdown in period $t+1$ if off-path quotes lead customers to ignore what they learned from past trades. Second, monotonicity is a natural refinement in the context of financial markets: Beliefs are consistent with the conviction that traders with higher estimates of the asset value offer higher prices, and would not want to offer prices that—if accepted—would lead to an immediate trading loss. Third, the unprejudiced-belief refinement implies that, in the pre-trade transparent market, dealers coordinate on the complete-information equilibrium once they share the same beliefs about the asset value. Without the refinement, it is possible to sustain "collusive" equilibria in which symmetrically informed dealers coordinate on high prices above the asset value, sustained by off-path beliefs of the customer who interprets any undercutting attempt as a signal of low value of the asset.

2.4 Equilibrium existence

Throughout, we assume that $q < 2\Delta^2/(1+\Delta^2)$ and

$$\delta\gamma \leq \bar{\delta}_{\text{exist}} \tag{2.1}$$

where $\bar{\delta}_{\text{exist}} > 0$ is a constant that we explicitly characterize in Lemma 2 in Appendix A.2. These assumptions guarantee equilibrium existence in all the regimes that we study. Intu-

[17] Note that these properties must hold automatically for on-path quotes b. The first part of the refinement is relatively standard in the literature; it is used in other dynamic signaling games such as Daley and Green (2012), Gul and Pesendorfer (2012), and Halac and Kremer (2020); both parts are similar in spirit to (but weaker than) the refinements of Cho and Kreps (1987) and Banks and Sobel (1987). The second part applies only when there exists a feasible dealer type that does not trade at a loss at the off-path quote. When no such type exists—for example, after a history that fully revealed the asset value, any quote below that value entails a trading loss for every feasible type—the restriction is silent.

[18] This refinement, which was introduced by Bagwell and Ramey (1991), and whose theoretical properties are studied by Vida and Honryo (2021), is commonly used in signaling games with multiple senders. It is also close in spirit to the refinement proposed by Milgrom and Mollner (2021).

itively, when both the persistence γ in the asset value and dealers' patience δ are high, a separating equilibrium may fail to exist when dealers' temptation to deviate by manipulating the market's future perception of the asset value becomes too strong. Additionally, the first of these conditions ensures that the amount of adverse selection is not too large relative to Δ—the private-value "buffer" between dealers and customers; otherwise, in regimes in which the customer may receive two quotes, it may be impossible to support a separating equilibrium in which the uninformed dealer breaks even.

Our modeling assumptions combined with the above restriction on parameters imply that our analysis is relevant for OTC markets in which the amount of private information about the asset value is not excessive. Moreover, the asset value changes relatively frequently, so that informational advantage of dealers is short-lived. This perspective complements papers like Kakhbod and Song (2020) and Kaya and Roy (2024) that study fully persistent private information and hence focus on the incentives of dealers to build reputations over time.

2.5 Welfare

Since the gain from trade is always Δ, welfare is proportional to the expected probability of trade. Let $\underline{\Lambda}_r \in [0, 1]$ and $\bar{\Lambda}_r \in [0, 1]$ denote the expected probability of trade, averaged and time-discounted over all periods, in regime r under the worst and best equilibrium, respectively, where $r \in \{\text{opaq, post, pre, full}\}$ indexes the four market regimes that we introduced in Section 2.2. We will rely on the following notions of welfare dominance.

Definition 1. Regime r weakly welfare-dominates regime r' if $\underline{\Lambda}_r \geq \underline{\Lambda}_{r'}$ and $\bar{\Lambda}_r \geq \bar{\Lambda}_{r'}$, with at least one inequality being strict. Regime r strongly welfare-dominates regime r' if $\underline{\Lambda}_r \geq \bar{\Lambda}_{r'}$ and $\bar{\Lambda}_r > \underline{\Lambda}_{r'}$.

Weak dominance of regime r over regime r' means that regime r leads to higher welfare both under the worst-case and the best-case equilibrium selection—the entire set of equilibrium payoffs shifts upwards under r.[19] However, it cannot be guaranteed that regime r will *always* lead to a better outcome than regime r'. Strong dominance is a more demanding but fully robust notion: If regime r strongly welfare-dominates regime r', then *any* equilibrium under regime r outperforms even the best equilibrium under regime r'.

[19]The reason why we only focus on the worst and the best equilibrium is that any level of welfare between these two extremes can also be achieved as long as a public randomization device is available.

3 Welfare comparisons of trading regimes

3.1 The opaque market

We begin our analysis with opaque markets, in which customers obtain a quote from a single dealer and do not have access to past transaction data.

3.1.1 Structure of equilibria

A full characterization of equilibria in the opaque market is provided in Proposition 1 in Appendix A.3 but we sketch the construction below. It is instructive to first consider a static monopoly benchmark in which there is only one dealer and trade happens only once. Recall that, by symmetry, we can focus on the case when the current customer in the market is attempting to *buy* the asset.

In a separating equilibrium of the static game, a quote b_σ of type $\sigma \in \{l, u, h\}$ reveals the dealer's information about the value of the asset to the customer; at the same time, the dealer with type σ must find it optimal to quote b_σ. Since in equilibrium we must have $b_l < b_u < b_h$, it follows immediately that types u and h cannot trade with probability one, as otherwise the low-type dealer would always have a profitable deviation. If types u and h trade at all, it must therefore be the case that $b_u = \Delta$ and $b_h = 1+\Delta$ since only these quotes make the consumer indifferent between accepting and rejecting, thus making it possible to implement an interior probability of trade by having the customer mix in equilibrium.

The quote of the low type is also pinned down: The low type quotes $b_l = -1+\Delta$, and the quote is accepted with probability one. To see that, note that since the customer's expected value for the asset is at least $-1+\Delta$ (regardless of beliefs), the low type can guarantee a profit of $\Delta - \epsilon$, for any $\epsilon > 0$, by quoting $-1+\Delta-\epsilon$ (which the customer strictly prefers to accept). On the other hand, a quote $b_l > -1+\Delta$ would be rejected by the customer. It follows that the low type must make a profit of Δ in equilibrium with a quote $b_l \leq -1+\Delta$, which is only possible when $b_l = -1+\Delta$ and b_l is accepted with probability one by the customer.

In equilibrium, no dealer-type can profitably deviate by mimicking the behavior of a different type, which leads to the following incentive constraints for the low and uninformed types:

$$\Delta \geq \lambda_u(\Delta - (-1)), \tag{IC$_{l\to u}$}$$

$$\lambda_u \Delta \geq \lambda_h((1+\Delta) - 0), \tag{IC$_{u\to h}$}$$

where λ_u and λ_h are the customer's equilibrium probabilities of accepting the quotes $b_u = \Delta$

and $b_h = 1 + \Delta$, respectively. It is easy to show that any (λ_u, λ_h) satisfying these conditions corresponds to an equilibrium, after specifying off-path beliefs of the customer appropriately. For example, we can set off-path beliefs to be the most pessimistic about the expected value of the asset subject to maintaining the monotonicity of the beliefs in the quote.[20] Moreover, any equilibrium is payoff-equivalent to some equilibrium of the form described above.

We refer to the above equilibrium construction as the *static monopoly trading pattern*: The low type quotes $-1 + \Delta$ which is accepted with probability one, the uninformed type quotes Δ which is accepted with probability $\lambda_u \in [0, \bar{\lambda}_u]$, and the high type quotes $1 + \Delta$ which is accepted with probability $\lambda_h \in [0, \bar{\lambda}_h]$, where the upper bounds on trade probabilities $\bar{\lambda}_u < 1$ and $\bar{\lambda}_h < 1$ can be explicitly calculated from the incentive constraints ($\text{IC}_{l\to u}$) and ($\text{IC}_{u\to h}$).

Even though the actual game we study is dynamic, post- and pre-trade opacity together imply that there is no learning between trading periods. Thus, under our stationarity assumptions, all equilibria consist of a repetition of the static monopoly trading pattern in every period. The most efficient equilibrium is characterized by binding incentive constraints ($\text{IC}_{l\to u}$) and ($\text{IC}_{u\to h}$), which pin down the highest probabilities of trade between a customer and the uninformed and high-type dealer, respectively, that can be sustained in equilibrium. These probabilities are bounded away from one. In the least efficient equilibrium, the uninformed and high-type dealer do not trade at all with a customer.

All equilibria are inefficient due to adverse selection: The customer interprets any attempt by the dealer to lower the quote as a signal of lower quality of the asset. The inefficiency is manifested through the randomized acceptance decision by the customer. Customers' mixing is a robust feature of all the equilibria we construct in the remainder of this section; it is important to emphasize that it is not a technicality—rather, it is shaped by fundamental economic forces. Due to adverse selection and signaling, in equilibrium, dealers post unattractive prices that make the customer indifferent between trading or not. Then, the customer accepts with probabilities that make dealers' pricing decisions optimal.[21] In the remainder of this section, we ask whether and to what extent various forms of trade transparency alleviate these informational frictions and lead to higher realized gains from trade.

[20]Formally, the customer believes that the expected asset value is 1 when the quote $b \geq 1 + \Delta$, 0 when $b \in [\Delta, 1 + \Delta)$, and -1 if $b < \Delta$. More generally, any system of off-path beliefs that attributes a quote $b < b_\sigma$ to a type strictly lower than σ with sufficiently high probability (provided that there is a lower type) supports the equilibrium.

[21]In Appendix OA.2, we consider an extension of our model in which period-t customer's private value for trading is $\Delta + \epsilon_t$, where ϵ_t is a random variable whose realization is privately observed by the customer. Then, acceptance decisions are deterministic, pinned down by a cutoff in the space of realizations of ϵ_t. Yet, dealers do not adopt pricing strategies that always lead to a transaction, due to adverse selection. Our simple model corresponds to the limiting case in which the variance of ϵ_t converges to zero.

3.2 The post-trade transparent market

We next turn to the post-trade transparent market. The difference to the opaque market is that any transaction becomes publicly disclosed to all future market participants.

3.2.1 Structure of equilibria

Before additional information arrives to the market through disclosure of past trades, equilibria exhibit a similar pattern as in the opaque regime. However, if the state is persistent in period t, the customer and the dealers learn from past trade disclosure, which affects equilibrium strategies. Specifically, if an informed dealer transacted in period $t-1$, equilibrium in period t becomes fully efficient. This is because the value of the asset v_t becomes common knowledge, dealers price the asset at $v_t + \Delta$, and the customer always accepts. This case captures the intuitive appeal of post-trade transparency: The informational content of prices is publicly revealed, and market participants become symmetrically informed. Thus, trade also becomes fully efficient, even when there is no direct competition between the dealers.

However, post-trade transparency does not alter the continuation equilibrium in case the outcome in period $t-1$ was either a trade by an uninformed dealer or no trade.[22] This is because the customer's beliefs still have full support over possible dealer types, and hence the adverse selection problem continues to constrain the equilibrium trade probabilities.

These observations are formalized in Proposition 2 in Appendix A.3 which fully characterizes equilibria of the post-trade transparent market.

Even though trade is efficient after disclosure of an informative transaction, dealers still extract all the surplus. Thus, "leveling the playing field" through post-trade transparency alone does not benefit the customers in otherwise opaque markets.

3.2.2 Is post-trade transparency beneficial?

Our first main result establishes a welfare ranking between full opacity and the post-trade transparent regime.

Theorem 1. *The post-trade transparent market weakly welfare-dominates the opaque market. The post-trade transparent market never strongly welfare-dominates the opaque market.*

The finding that post-trade transparency leads to a welfare improvement may seem intuitive given the fully efficient outcome following an informative trade. However, the result

[22] Recall that—even though we focus on the case that the customer is a buyer in describing the equilibrium—the customer is a seller with equal probability. This is why the event of no trade is equally likely to occur under both high and low asset value.

is not *a priori* obvious because introducing post-trade transparency also changes dealers' incentives when quoting in periods in which the state resets. On the one hand, post-trade transparency creates additional incentives for dealers to mimic a higher type when $Z_t = 0$ (under our maintained convention that the customer is a buyer). By quoting a higher price, dealers can now manipulate the beliefs of period-$t+1$ customers whenever their quote is executed. This tends to *tighten* incentive constraints. On the other hand, by ensuring efficient trading at the monopoly price in the subsequent period, post-trade transparency increases dealers' on-path continuation payoffs, which tends to *relax* incentive constraints. Under both the worst and the best equilibrium, the latter effect dominates the former.

In the worst equilibrium, the uninformed and the high type do not trade at $Z_t = 0$. As a result, an upward manipulation of market beliefs is simply not possible because a deviation to a higher quote results in no transaction (and is hence not observed by future market participants). Therefore, the main drawback of post-trade transparency is vacuous in the worst equilibrium. At the same time, the main positive aspect of post-trade transparency—efficient trade after information becomes symmetric—materializes whenever the low-type dealer trades and the state is persistent. Thus, welfare in the worst-case equilibrium is strictly higher with post-trade transparency.

In the best equilibrium, post-trade transparency improves informed dealers' on-path equilibrium continuation payoffs conditional on trading. Indeed, if an informed dealer trades in period $t-1$, then she trades in period t at the monopoly price with probability one (conditional on being contacted and the state being persistent but regardless of the type of the customer). A deviation to a higher price at $t-1$ always lowers the probability of trading, meaning that informed dealers face a higher opportunity cost of not trading in the post-trade transparent regime. This positive effect relaxes the incentive constraint associated with a potential deviation of a low-type dealer to the uninformed dealer's quote: The low type's equilibrium payoff is higher under post-trade transparency while the deviation is equally attractive as in the opaque market (since the uninformed quote does not lead to an update in market beliefs). Because deviating to an uninformed quote becomes less attractive for the low type, the uninformed dealer can trade with higher probability in equilibrium. At least in the best equilibrium, this additional positive effect is what overturns the uninformed dealers' increased "temptation" to deviate to a high quote.

This logic suggests that the benefits of post-trade transparency are subtle. In equilibrium, informed dealers do not want to manipulate market beliefs in part because manipulation requires quoting higher prices that are often rejected by the customer, so the manipulator must forgo some of the benefits of fully efficient trade in the continuation game. Dealers benefit from efficient trade because they extract all the surplus from customers. In the earlier working paper Vairo and Dworczak (2024), we showed that the conclusion of Theorem 1 may sometimes fail when customers can contact dealers sequentially and search costs are low—the opaque market may have a more efficient best equilibrium than the post-trade transparent market for intermediate values of q. This is because dealers do not extract all the surplus from customers in equilibria with search, and therefore whether the positive or negative effects of post-trade transparency dominate is parameter-dependent.

Given the above discussion, it is perhaps unsurprising that the post-trade transparent market never *strongly* welfare-dominates the opaque market. This means that the best equilibrium in the opaque market is more efficient than the worst equilibrium in the post-trade transparent market. Thus, post-trade transparency regulation (such as TRACE) may lower welfare if market participants coordinate on an inefficient equilibrium.

3.2.3 Discussion and related literature

Post-trade transparency has received significant attention in the literature, so Theorem 1 is related to several existing results, albeit derived in different settings.

Duffie et al. (2017) analyze the transparency effects of introducing a benchmark (a form of post-trade transparency) into an opaque search market and find that benchmarks raise surplus when search costs are sufficiently large. Their model does not feature adverse selection, and the benefits of transparency operate primarily through an improved bargaining position for customers. A result more closely related to Theorem 1 is offered by Back et al. (2020), who show that post-trade transparency induces intermediaries to offer more attractive prices to sellers when these prices are observed by prospective buyers, due to a signaling motive. In their model, post-trade transparency increases upstream trade volume and hence efficiency; however, the benefits of signaling are strengthened by the assumption that downstream trade between an intermediary and a buyer always occurs. Ollar et al. (2021) analyze the disclosure of past trades in a dynamic uniform-price double auction and find that transparency generally increases efficiency but may need to be limited for equilibria to exist; this is echoed by our finding that separating equilibria can fail to exist when dealers can profitably manipulate beliefs in the market.

Belief manipulation is not the only potential cost of post-trade transparency. Asriyan et al. (2017) analyze a dynamic model with adverse selection and spillovers across markets. They show that sellers may strategically delay trade in anticipation of information arriving from markets for correlated assets. Similar to us, they find equilibrium multiplicity that leads to ambiguous effects of disclosing past trades on welfare. Kaya and Roy (2024) study disclosure of past transaction volume in a dynamic model in which uninformed buyers make offers to a seller who is privately informed about asset quality. They find that the effects of transparency can be positive or negative depending on (i) the market's initial belief about asset value and (ii) the intensity of competition between buyers, due to a subtle impact on the seller's incentives to build a reputation for high quality. Kaya and Roy (2022b) further show that observability of past prices may improve market efficiency in some cases.[23]

Theorem 1 is broadly consistent with these results: post-trade transparency can increase efficiency by symmetrizing information but subject to several caveats. Welfare improvement is not fully robust to equilibrium selection; the possibility of belief manipulation limits the benefits of disclosure; and trade remains highly inefficient when past transactions are uninformative about value. Next, we examine whether pre-trade transparency addresses these concerns more effectively.

3.3 The pre-trade transparent market

In this section, we consider the trade-off between post- and pre-trade transparency. Effectively, we ask whether it is better for market participants to observe past executed quotes or quotes currently available in the market.

3.3.1 Structure of equilibria

It is instructive to first consider how pre-trade transparency affects trading in a static model, in which the customer observes both quotes simultaneously and then decides with whom (if at all) to trade—inducing a first-price auction between the dealers.

For $\sigma \in \{l, u, h\}$, let b_σ be a quote in the support of type σ's equilibrium strategy. Since the high-type dealer can only win a trade against another high type in equilibrium, the price-undercutting logic (the possibility of deviation to a price ϵ below the competitor's price) implies that $b_h = 1$ with probability one. A similar argument establishes that the uninformed dealer must play a pure strategy with $b_u \leq \Delta$. If $b_u < \Delta$, then the customer finds it optimal to accept after receiving a pair of quotes (b_u, b_u). If $b_u = \Delta$, then the customer is indifferent between accepting (either quote) and rejecting (both quotes), and we use $\lambda_u \in [0, 1]$ to denote the equilibrium probability that the customer trades after seeing (Δ, Δ). Finally, a standard argument implies that the low type's quoting strategy is atomless with support on an interval $[\underline{b}_l, -1 + \Delta]$ with $\underline{b}_l > -1$. The continuous distribution of quotes for the low type results from the quantity-price trade-off: Each quote in the support wins against an uninformed competitor, but lower quotes have a higher chance of winning against a low-type competitor.

[23]See also Kaya and Roy (2022a), who show that gains from trade can be non-monotonic in the length of the record of past trades.

This argument establishes that there are two types of candidate equilibria. In the first type of candidate equilibrium, the uninformed dealer quotes $b_u < \Delta$ and the outcome is fully efficient. However, this equilibrium is sustainable only if the low type does not benefit from deviating to b_u. Because the low-type dealer's strategy is atomless, the probability of trading against a low-type competitor when she quotes $-1+\Delta$—which is the supremum of the support of her equilibrium strategy—is zero. Therefore, her equilibrium payoff is $(1-q)\Delta$. If she deviates to b_u she only trades when the competitor is uninformed and the tie is broken in her favor. Thus, her incentive constraint is

$$(1-q)\Delta \geq \frac{1-q}{2}(b_u+1).$$

Given that $b_u \geq 0$ is required for the uninformed dealer to at least break even, an efficient equilibrium does not exist when $\Delta < 1/2$. We derive an if-and-only-if condition for the existence of the efficient equilibrium in Appendix A.4.

In the second type of equilibrium, in which strategies feature $b_u = \Delta$, trade is inefficient when both dealers are uninformed. The values of λ_u that can arise in equilibrium are determined by the incentive-compatibility constraints preventing the low and the uninformed types from mimicking the adjacent higher type:

$$(1-q)\Delta \geq \frac{(1-q)\lambda_u}{2}(1+\Delta), \tag{IC$_{l\to u}$}$$

$$\frac{(1-q)\lambda_u}{2}\Delta + \frac{q}{2}(\Delta-1) \geq 0. \tag{IC$_{u\to h}$}$$

We formalize these arguments in the proof of Proposition 3 in Appendix A.4, where we show that any λ_u satisfying conditions (IC$_{l\to u}$) and (IC$_{u\to h}$) can be part of an equilibrium. Off-path beliefs need to be specified only when exactly one quote is on path (the case when both quotes are off-path is irrelevant). In the construction, the customer's belief is always consistent with the on-path quote; but if the on-path quote is uninformative (i.e., equal to b_u), the customer updates beliefs based on the off-path quote in the same way as in the opaque market.

The *simultaneity* of quotes received by the customer is important. Suppose that high-type dealers quote a price $b_h > 1$. This cannot be an equilibrium in the pre-trade transparent regime because a high-type dealer has a profitable deviation to quote $b_h - \epsilon$ for small enough $\epsilon > 0$. This lower off-path quote must be accepted by the customer with probability one (in the relevant contingency when both dealers have a high type) because the customer learns that the value of the asset is high from observing the other dealer's quote equal to b_h.[24]

[24]This off-path updating rule—despite its intuitive appeal—does not follow from the concept of perfect

However, consider a model in which the customer observes the two quotes sequentially. In that case, it is possible to support $b_h > 1$ in equilibrium, and trade can fail if $b_h = 1 + \Delta$. This is because the customer may believe that an off-path quote $b_h - \epsilon$ signals a lower asset value and therefore reject it. This observation highlights the fragility of competitive forces under asymmetric information: Ensuring that customers can easily access multiple quotes is not, by itself, sufficient to induce competition among dealers.

3.3.2 Which is better: post- or pre-trade transparency?

The discussion of the equilibrium structure may suggest a trade-off between post- and pre-trade transparency. On the one hand, pre-trade transparency symmetrizes information across dealers immediately after the arrival of new information, which should improve efficiency in periods t with $Z_t = 0$. On the other hand, post-trade transparency induces fully efficient trade in periods t with $Z_t = 1$ following an informative transaction in period $t-1$. It turns out, however, that pre-trade transparency is always the superior market design. In particular, it is also more efficient in periods with $Z_t = 1$. This is because pre-trade transparency symmetrizes information across dealers regardless of whether the transaction in period $t-1$ was informative about the asset value, whereas post-trade transparency does so only following an informative trade.

Theorem 2. *Pre-trade transparency weakly welfare-dominates post-trade transparency. When q is sufficiently high, pre-trade transparency strongly welfare-dominates post-trade transparency.*

The proof of Theorem 2 shows that pre-trade transparency is strictly more efficient than post-trade transparency both in the worst and in the best equilibrium *conditional on every realization of Z_t.*

In periods in which the value of the asset is redrawn ($Z_t = 0$), more competition in the pre-trade transparent regime induces dealers to quote lower prices. This sustains a higher probability of trade and raises welfare. In particular, trade occurs with certainty whenever at least one dealer is informed.[25] In the post-trade transparent regime, this is not true: In all equilibria, a customer who contacts an uninformed dealer must exit the market with positive probability (regardless of the type of the other dealer).[26] Intuitively, the uninformed type cannot trade for sure under pre-trade opacity because that would create a profitable

Bayesian equilibrium alone, which is why we imposed the unprejudiced beliefs refinement.

[25] If the probability q of an informed type is high enough, this makes the pre-trade transparent market strongly welfare dominate the post-trade transparent market.

[26] As we showed in the earlier working paper Vairo and Dworczak (2024), this conclusion is true even when the customer can search at an arbitrarily low cost.

deviation for the low-type dealer. With pre-trade transparency, however, a low-type dealer is disciplined by the possibility of competing against another low type, since an upward price deviation leads to losing the trade in that scenario. As a result, uninformed (and hence high) types can trade with higher probability in equilibrium.

The pre-trade transparent market also dominates the post-trade transparent market when the state is persistent ($Z_t = 1$). The post-trade transparent market is fully efficient after an informative transaction at $t-1$ but is not fully efficient otherwise. In contrast, the pre-trade transparent regime *always* features efficient trade at $Z_t = 1$. The reason is that the dealers' information becomes symmetric following the period-($t-1$) auction, regardless of its outcome, since quotes in the auction are assumed to be observed by dealers. Symmetric information, combined with direct competition, is sufficient to ensure efficiency, even if customers remain uninformed about the value of the asset.

To illustrate this point, consider the case when an uninformed and a high-type dealer are present in the market at $Z_{t-1} = 0$. With post-trade transparency, the efficient outcome is implemented at $Z_t = 1$ only if the high type traded at $Z_{t-1} = 0$. In all other cases, disclosure does not refine the customer's beliefs sufficiently to guarantee efficiency since she will still be concerned about adverse selection when evaluating quotes from any given dealer. With pre-trade transparency, however, it becomes common knowledge across *dealers* (after observing each other's quotes in period $t-1$) that the value of the asset is high. Even though the customer in period t is completely uninformed, direct competition between symmetrically informed dealers results in a competitive offer equal to the true value of the asset, resulting in a fully efficient outcome.

3.3.3 Discussion and related literature

The allocative efficiency properties of auctions are well understood in economics (going at least as far back as Vickrey, 1961). Since the pre-trade transparent market is effectively a repeated first price auction, Theorem 2 may seem ex ante unsurprising. However, the presence of adverse selection in the model makes this result far less obvious.

First, simultaneity of quotes is crucial to ensure that competitive forces kick in. It would be tempting to conjecture that a pre-trade transparent (auction) market should have similar welfare to a search market with vanishing search costs. After all, the customer can obtain the same information in both regimes at an arbitrarily low cost in the latter. Barring the Diamond-paradox outcome (Diamond, 1971)—e.g., by assuming that a small fraction of customers visit all dealers as in Janssen et al. (2011) or Duffie et al. (2017)—this would indeed be the case if dealers' information pertained only to their private values for the asset. As we have argued, however, this intuition is misleading when dealers have private information relevant to the customer's evaluation of the asset. In such settings, inefficiently high prices can be sustained in equilibrium by adversarial inferences from off-path quotes. Lester et al. (2019) make a related point, showing that search markets may be highly inefficient in the presence of adverse selection.

Second, and more novel to our framework, it is important that dealers observe *each other's* quotes. This observation adds an important caveat to Theorem 2: The unambiguous welfare comparison would not hold in general if the auction format did not allow dealers to see each other's quotes before the next trading period. This point is relevant to the debate on disclosure of unexecuted quotes on trading platforms. In practice, trading platforms adopt different rules governing the disclosure of the "cover"—the runner-up offer—and sometimes give the traders discretion over that decision (SIFMA, 2016). Ordin and Sverchkov (2024) study the welfare effects and the incentives of customers to disclose rejected quotes to the winning dealer in a model of repeated trading under adverse selection. They show that the customer may sometimes choose to conceal that information. Our results predict that disclosure of all quotes (including the cover) to all participating dealers is beneficial for trade efficiency. If we assumed that rejected quotes were unobservable, dealers would only be able to partially infer their competitor's type.[27] As a result, information asymmetry would persist at $Z_t = 1$ and a fully efficient outcome would no longer be guaranteed.

When the likelihood of informed dealers is high, Theorem 2 establishes a strong conclusion, which is that even the worst equilibrium of the pre-trade transparent market is more efficient than the best equilibrium of the post-trade transparent market. This provides a particularly robust argument in favor of auctions in markets where dealers have a significant informational advantage over customers.

We are not aware of any prior work directly comparing post- and pre-trade transparency. Glebkin et al. (2023) look at the effects of introducing an auction between multiple traders in the canonical OTC model of Duffie et al. (2005). Their analysis does not feature adverse selection and focuses on the effects on the equilibrium distribution of asset holdings in a large market. Riggs et al. (2020) find that market imperfections, such as private information and existing relationships between dealers and customers, may prevent pre-trade transparent regimes from being fully efficient. They study a model in which dealers are exposed to a form of winner's curse that—unlike in our paper—manifests in dealers' future (in)ability to rebalance their positions in the inter-dealer market. Finally, as explained earlier, Kaya and Roy (2024) study the interaction of post-trade transparency and competition, but since

[27] From a technical point of view, this would significantly complicate our analysis; because dealers' inference about their competitor's type would now depend on their quote (and the low type quotes a price from a continuous distribution), there would be a continuum of different continuation equilibria in period $Z_t = 1$ differing in the dealers' beliefs.

the competition is between uninformed parties in their framework, there is no within-period quote dispersion in the market, and hence pre-trade transparency is a vacuous concept.

3.4 The fully transparent market

Finally, we turn attention to the fully transparent market in which both the current quotes and the past transaction prices are observed.

3.4.1 Structure of equilibria

The structure of equilibria in the fully transparent market is the same as in the pre-trade transparent market (see Proposition 4 in Appendix A.5 for a full equilibrium characterization). The key difference is that, under full transparency, future customers also observe transaction prices. This enables dealers to engage in belief manipulation, which modifies the equilibrium probabilities of trade for the uninformed type.

To see why that is the case, consider a low-type dealer contemplating a deviation to the quote Δ of the uninformed dealer at $Z_t = 0$. Suppose that the other dealer is uninformed and hence also quotes Δ. Without post-trade transparency, the informational effect of this deviation (relative to the equilibrium path) is that the uninformed competitor will not learn that the value of the asset is low. With post-trade transparency, this deviation will also manipulate the future customer's belief: Having observed either no trade or a transaction with price Δ, the future customer will conclude that neither dealer can have a low type. No longer concerned about potential adverse selection, the future customer will agree to trade at any price below Δ. This allows the deviating low-type dealer to make additional profits in period $t + 1$ by undercutting the uninformed competitor—whose continuation strategy is to quote 0 if $Z_{t+1} = 1$. This deviation is attractive because the dealer knows that the actual value of the asset is low.

In contrast, without post-trade transparency, the future customer will have full-support beliefs about dealers' types. In equilibrium, she may decline to trade at off-path quotes below 0, interpreting them as coming from the low type.[28] Thus, the deviating low-type dealer can no longer undercut the competitor in period $t + 1$—she must either match the competitor's quote (in which case she only trades with probability 1/2) or lower prices more aggressively to $-1 + \Delta$ to guarantee that the customer will be willing to trade regardless of her beliefs. In either case, the deviation is less profitable without post-trade transparency. Hence, the additional disclosure of transaction prices under full transparency strengthens the low type's

[28]In particular, this will happen in the most efficient equilibrium, in which off-path beliefs are set to minimize the payoffs from deviations.

incentive to mimic the uninformed type, tightening the relevant incentive constraint without generating any offsetting efficiency gains.

3.4.2 Is full transparency beneficial?

Our next result formally compares the pre-trade transparent and fully transparent markets.

Theorem 3. *There exists $\bar{\delta}_{\text{full}}$ such that, if $\delta\gamma \leq \bar{\delta}_{\text{full}}$, pre-trade transparency and full transparency have the same set of equilibrium levels of welfare. Otherwise, pre-trade transparency weakly welfare-dominates full transparency.*

The main conclusion of Theorem 3 is that adding post-trade transparency on top of pre-trade transparency does not improve—and may actually harm—market efficiency. There are two pieces of intuition for this result. First, post-trade transparency does not help in any contingency. Since dealers are already symmetrically informed after one round of trading due to pre-trade transparency, the only effect of the public announcement of past trades is that future customers also become informed about the value of the asset. However, as we emphasized previously, ensuring that customers are informed is not actually needed for market efficiency. Second, post-trade transparency may harm efficiency because it tightens the incentive-compatibility constraints for the dealers who can now benefit from manipulating future customers' beliefs, as explained above.

Summarizing the argument, full transparency enables informed dealers to engage in more effective manipulation by jointly deceiving an uninformed competitor *and* the future customer. With pre-trade transparency alone, manipulation does not affect the future customer's beliefs and—at least in the most efficient equilibrium—it is more difficult for the informed dealer to exploit the informational advantage over the competitor. Since the belief-manipulating deviation is more attractive for the informed dealer under full transparency, the probability of trade following a pair of uninformed quotes must be lower to make the deviation unprofitable. Thus, the most efficient equilibrium exhibits lower probability of trade when post-trade transparency is added to the market.

3.4.3 Discussion and related literature

To the best of our knowledge, the idea that post-trade transparency may provide no additional benefits when layered onto a pre-trade transparent market is new to the literature. That said, the fact that belief manipulation can constrain the benefits of disclosure has been examined in other contexts (see, e.g., Dworczak, 2020; Antill and Duffie, 2020; Zhang, 2022 for recent examples).

While Theorem 3 offers an unambiguous welfare comparison, we emphasize that the superiority of pre-trade transparency requires several qualifications. First, the welfare ranking is weak: The worst-case equilibrium is the same across the two regimes, while the best-case equilibrium is weakly more efficient under pre-trade transparency (strictly, unless both regimes have a fully efficient equilibrium). Thus, introducing post-trade transparency may increase welfare if it induces market participants to coordinate on a more efficient equilibrium.

Second, the benefits of post-trade transparency are muted because pre-trade transparency alone leads to fully efficient outcomes when the state is persistent. As emphasized in the previous section, full efficiency at $Z_t = 1$ hinges on dealers observing each other's quotes in the pre-trade transparent market. If this were not the case, there would generally be a trade-off associated with introducing post-trade transparency.

Third, full efficiency in the pre-trade transparent market at $Z_t = 1$ could fail under weaker equilibrium refinements. In Appendix OA.3, we show that without refinements on off-path beliefs, there exist inefficient equilibria in which symmetrically informed dealers maintain prices above the true value of the asset, despite Bertrand-style competition. While we view these equilibria as unnatural, they could arise as a form of collusion between the dealers. In such equilibria, dealers exploit their *joint* informational advantage over the uninformed customer. Post-trade transparency can then help prevent these collusive outcomes by making customers more informed. In the appendix, we show that—without our belief refinements—worst-case equilibria become more efficient when post-trade transparency is added. At the same time, however, relaxing the refinements makes the best-case equilibria in the pre-trade transparent regime even more efficient. Therefore, allowing for these collusive equilibria could make the ranking presented in Theorem 3 ambiguous, but it would not revert it.

4 Optimal trading mechanism

In this section, we use mechanism-design tools to construct a trading mechanism that maximizes welfare. The main objective is to test the robustness of our predictions from Section 3. Studying direct mechanisms allows us to abstract away from some of the limiting assumptions on parameter values, the nature of equilibria, and off-path beliefs that we made in the equilibrium analysis.

We use a relatively stringent notion of implementability—less permissive than would be implied by the solution concept used in Section 3—and show that the optimal mechanism nevertheless outperforms all of the regimes analyzed so far. This is in part because the mechanism-design framework allows the market designer to change more than just market transparency. Nevertheless, the exercise is useful because it reveals which intuitions and results from the previous section continue to hold at the optimal mechanism. Moreover, we show that the optimal outcome can be implemented by appropriately modifying the pre-trade transparent regime (without any post-trade transparency), thereby reinforcing the main takeaways from Section 3.

By the revelation principle, we can restrict attention to direct mechanisms in which agents report their exogenous private information to the mechanism, and the mechanism specifies the trading probabilities and transfers. We study ex-post implementation, requiring that agents' reports and participation decisions remain optimal even conditional on learning other agents' types.[29] Ex-post implementation rules out the fully-efficient but economically unappealing mechanisms of Crémer and McLean (1985). Finally, to parallel the stationary structure of equilibria we constructed in Section 3, we assume that the mechanism can only condition on information that is relevant for the current asset value.[30]

To set up the problem formally, we let $\theta_{i,t} \in \Theta \equiv \{-1, 0, 1\}$ denote dealer i's type in period t with $Z_t = 0$. For notational convenience, we redefined the types so that they are equal to the corresponding expected value of the asset given the exogenous signal the dealer receives about v_t. Given that dealers do not receive any exogenous information in periods in which the state is persistent, we set $\theta_{i,t} = \theta_{i,t-1}$ when $Z_t = 1$. We denote $v(\theta_i, \theta_{-i}) \equiv \mathbb{E}[v_t|\theta_{i,t} = \theta_i, \theta_{-i,t} = \theta_{-i}]$. Since the state Z_t was assumed common knowledge, we assume that it is also observable to the designer (as will be clear from our results, this assumption is not needed to implement the optimal mechanism). A fully general dynamic mechanism conditions outcomes at time t on the entire history of reports and states up to and including time t. We instead restrict the designer to use a *stationary mechanism* that conditions outcomes at time t only on period-t reports, the realization of Z_t, and reports in $t-1$ if $Z_t = 1$. Because the only players who possess private information are dealers at times t when $Z_t = 0$, we can further restrict attention to mechanisms in which dealers report their types to the mechanism only when $Z_t = 0$.

Overall, given that a symmetric mechanism is without loss of generality in our setting, we define a trading mechanism to be $(x_Z, p_Z, t_Z) : \Theta^2 \to [0,1] \times \mathbb{R}^2$, for $Z \in \{0, 1\}$, where $x_Z(\theta_{i,t}, \theta_{-i,t})$ is dealer i's probability of trading the asset at time t when $Z_t = Z$, i reported $\theta_{i,t}$, and $-i$ reported $\theta_{-i,t}$ (in the period in which information arrived); $p_Z(\theta_{i,t}, \theta_{-i,t})$ is the

[29] Unlike in Section 3, we do not impose any belief refinements because the direct-mechanism analysis does not involve off-path beliefs. We also use ex-post implementation in the indirect implementation in Theorem 5 but we consider belief refinements in the unique implementation result in Appendix OA.4.

[30] This rules out mechanisms that can punish a dealer for misreporting their type by changing the future trading protocol—a feature that not only would be ruled out by the equilibrium concept from Section 3 but also seems unrealistic.

corresponding monetary transfer that dealer i receives, and $t_Z(\theta_{i,t}, \theta_{-i,t})$ is the transfer paid by the customer. To simplify exposition, we again restrict attention to the case when the customer is a buyer.[31] A mechanism is *feasible* if it satisfies:

- *Dealers' ex-post incentive compatibility:*

$$p_0(\theta_i, \theta_{-i}) - x_0(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i}) + \delta\gamma[p_1(\theta_i, \theta_{-i}) - x_1(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i})] \geq$$
$$p_0(\theta_i', \theta_{-i}) - x_0(\theta_i', \theta_{-i})v(\theta_i, \theta_{-i}) + \delta\gamma[p_1(\theta_i', \theta_{-i}) - x_1(\theta_i', \theta_{-i})v(\theta_i, \theta_{-i})], \quad \forall \theta_i, \theta_i', \theta_{-i} \in \Theta;$$

- *Dealers' ex-post per-period individual rationality:*

$$p_Z(\theta_i, \theta_{-i}) - x_Z(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i}) \geq 0, \quad \forall \theta_i, \theta_{-i} \in \Theta,\ Z \in \{0, 1\};$$

- *Customers' ex-post individual rationality:*

$$(v(\theta_A, \theta_B) + \Delta)(x_Z(\theta_A, \theta_B) + x_Z(\theta_B, \theta_A)) - t_Z(\theta_A, \theta_B) \geq 0, \quad \forall \theta_A, \theta_B \in \Theta,\ Z \in \{0, 1\};$$

- *Ex-post per-period budget balance for the designer:*

$$t_Z(\theta_A, \theta_B) \geq p_Z(\theta_A, \theta_B) + p_Z(\theta_B, \theta_A), \quad \forall \theta_A, \theta_B \in \Theta,\ Z \in \{0, 1\};$$

- *Unit-demand constraint:*

$$x_Z(\theta_i, \theta_{-i}) \in [0, 1] \text{ and } x_Z(\theta_i, \theta_{-i}) + x_Z(\theta_{-i}, \theta_i) \leq 1, \quad \forall \theta_i, \theta_{-i} \in \Theta, Z \in \{0, 1\}.$$

Ex-post incentive compatibility requires that it is optimal for dealers to truthfully report their type for every realization of customer's desired positions and the competitor's type, and conditional on the competitor reporting truthfully.[32] Ex-post individual rationality for dealers and customers means that market participants can walk away from a trade at any time if they find it unprofitable, which is a stronger assumption than the one made in Section 3 (where our solution concept only guaranteed this property for the customers).

[31] This restriction is without loss of generality. As we show in Appendix A.6, the only relevant dynamic incentive compatibility constraint for the dealers is the one where, at $Z_t = 0$, the dealer knows that the customers that she will face before the state resets are either both buyers or both sellers. As a result, there are no interactions between the customer-buyer and customer-seller cases in the designer's problem, which implies that the two problems can be solved separately and their solutions are symmetric.

[32] The restriction to stationary mechanisms implies that this condition is independent of continuation payoffs conditional on the state being redrawn in the future.

Ex-post budget balance states that the designer cannot subsidize trades in any period.[33] The unit-demand constraint simply ensures that trading probabilities add up to less than one.

An *optimal mechanism* is one that maximizes the total discounted probability of trade,

$$\sum_{t=0}^{\infty} \delta^t \, \mathbb{E}[x_{Z_t}(\theta_{i,t}, \theta_{-i,t}) + x_{Z_t}(\theta_{-i,t}, \theta_{i,t})],$$

over all feasible mechanisms, where the expectation is taken over the sequence $(\theta_{i,t}, \theta_{-i,t}, Z_t)_{t=0}^{\infty}$.

Our first result in this section provides a characterization of an optimal mechanism.

Theorem 4 (Optimal trading mechanism)**.** *There exists an optimal mechanism which is static: The implemented outcome is the same in each period conditional on dealers' reports. Expected welfare in an optimal mechanism is higher than in the best-case equilibrium of the pre-trade transparent market (strictly, unless the pre-trade transparent regime achieves the fully efficient outcome). In each period, the dealer with the lower type trades with the customer who pays a price equal to the other dealer's type (with ties broken uniformly at random), except for two cases:*

- *If both dealers are uninformed, trade happens with probability* $\min\{1, 2\Delta\}$*;*
- *If one dealer has a low type and the other dealer is uninformed, the customer pays a price equal to* $-1 + \Delta$*.*

The first conclusion from Theorem 4 is that there exists an optimal mechanism that takes a *static* form, independent of Z_t. To see why, note that under our restriction to stationary mechanisms, every time the state is redrawn the designer faces the same optimization problem. Since the state remains fixed until the next arrival of news, the problem naturally decomposes into an infinite sequence of identical two-period problems: one period immediately after a redraw (when $Z_t = 0$), followed by the continuation while the state persists (when $Z_t = 1$). Within each of these problems, dealers' types are therefore fully persistent. It follows from Baron and Besanko (1984) that the optimal dynamic mechanism in this two-period problem is simply the repetition of the optimal static mechanism.[34]

The second conclusion is that the optimal mechanism is fully efficient when $\Delta \geq 1/2$. Intuitively, when gains from trade are sufficiently large, the customer can compensate the

[33] This assumption is again made to parallel the setting of Section 3 but it is with loss of generality. We can show that, whenever $\Delta < 1/2$, the designer can strictly increase total welfare if she were allowed to run ex-post budget deficits while still balancing the budget on average.

[34] This is no longer true when the mechanism is allowed to be non-stationary. In that case, the designer can use continuation payoffs to relax dealers' incentive-compatibility constraints, so the optimization problem following a redraw is no longer static. For example, the designer could condition future tie-breaking rules on past reports whenever dealers' reports coincide.

low-type dealer enough to discourage her from mimicking the uninformed type while still obtaining nonnegative expected surplus from trade. Thus, the incentive constraints can be satisfied without sacrificing efficiency. This is in line with the intuition that, in the presence of adverse selection, efficient outcomes may become attainable if the gains from trade are sufficiently large. However, it is worth highlighting that the condition $\Delta \geq 1/2$ is not enough to ensure efficient equilibrium trading in any of the trading regimes that we discussed before. When $\Delta < 1/2$, efficient trade is no longer implementable. To sustain trade with probability one when both dealers are uninformed, the mechanism must provide the low type with sufficiently high transfers to deter her from mimicking the uninformed type. When gains from trade are too small, however, such transfers would force the customer to trade at a loss. Consequently, every incentive-compatible mechanism must sacrifice trade with strictly positive probability. Theorem 4 implies that this inefficiency is not a consequence of a suboptimal trading protocol but a fundamental information friction in our trading environment. At the same time, none of the regimes discussed in Section 3 addresses this friction in an optimal way.

The final conclusion from Theorem 4 is that the optimal mechanism resembles a second-price auction (SPA). In an ex-post equilibrium of the SPA, trade happens with the lowest-type dealer and the price is equal to the valuation of the losing dealer. The mechanism in Theorem 4 departs from the outcome of the SPA in two ways: Trade happens with interior probability when both dealers are uninformed and $\Delta < 1/2$, and the price lies strictly in between the winning and the losing quote when one dealer has a low type and the other dealer is uninformed. The first of these departures is a consequence of informational asymmetries between dealers, as described above. The second is a consequence of the assumption that the buyer must have ex-post incentives to participate in the trade. In particular, conditional on the profile of quotes coming from types $(\theta_A, \theta_B) = (-1, 0)$, the customer would be willing to trade at the winning quote but not at the losing quote (since she learns that the value of the asset is low). Hence, the optimal mechanism modifies the price to ensure customer participation.

4.1 An optimal time-invariant trading protocol

Theorem 4 is silent about market transparency. A consequence of the revelation principle for dynamic interactions (see, e.g., Myerson, 1982) is that optimality never requires the designer to disclose information to the agents, other than by recommending actions they should take (here, trade outcomes). Intuitively, transparency never helps satisfy incentive-compatibility or individual-rationality constraints because disclosure creates more contingencies in which these constraints must hold—the designer can do weakly better by "pooling incentives." However, in reality, traders are able to learn from the observed outcomes of each trading round, and our focus on ex-post equilibria implies that certain information can be disclosed to them without violating incentive compatibility or individual rationality.[35] These insights allow us to construct a realistic implementation of the optimal mechanism from Theorem 4 that casts light on the optimal degree of market transparency.

Consider the following market design, which we refer to as the *modified second-price auction*. Dealers simultaneously submit quotes; the dealer with the lower quote wins but the trading price is equal to the higher quote. Ties are broken uniformly at random. The customer observes dealers' behavior in the auction. After the auction, the customer decides between one of three actions: agree, exit, or ask for a new quote. If she agrees, she pays the auction price to the winning dealer and a fee equal to Δ to the platform. If she exits, there is no trade. If the customer requests a new quote, the winning dealer makes an additional take-it-or-leave-it offer, which the customer either accepts or rejects. If she accepts, trade takes place at this offer (with no additional fee); if she rejects, there is no trade. Dealers observe each other's quotes in the auction. There is no post-trade transparency.

Theorem 5 (Indirect implementation of the optimal mechanism)**.** *The allocation of the optimal mechanism described in Theorem 4 is an ex-post equilibrium outcome of the protocol in which the modified second-price auction is run in every period.*

The proof constructs equilibrium strategies of the modified SPA that implement the optimal outcome. When $Z_t = 0$, dealers' equilibrium strategy is to quote their estimates of the asset's value. The auction rules (and the fee structure) imply that the customer always pays her expected value for the asset conditional on the observed quotes and hence is willing to trade, except when one dealer has a low type and the other dealer is uninformed. In this event, the customer learns that the asset value is low and hence she is not willing to trade at the price 0 prescribed by the auction. However, she would agree to trade at the low type's quote of -1; thus, the customer requests a new quote from the dealer. Because the dealer has all the bargaining power at this stage of the game, she extracts all the surplus and trade happens with probability one at the price $-1 + \Delta$. When $Z_t = 1$, the equilibrium strategies retain a similar structure but are appropriately modified to account for the fact that dealers learned each other's types in the previous-period auction (while the period-t customer is uninformed due to post-trade opacity).

[35] In a dynamic game, by an ex-post equilibrium, we mean a perfect Bayesian equilibrium in which each player's behavior remains optimal on-path even if she learned the other players' private information—here, the dealers' types. Note that this does not dispense with off-path beliefs: After an observed deviation, the remaining players must still form beliefs about the deviator's type.

While all ingredients of our indirect mechanism—the SPA pricing rule, trading fees, and post-auction bargaining—are used on some existing OTC trading platforms (SIFMA, 2016), it is *not* our goal to argue that this particular design should be used in practice. The protocol from Theorem 5 inherits some of the stylized features of our framework (e.g., the fact that the fee is exactly Δ, or that post-auction bargaining is only needed in one contingency) that would not be present in a richer environment. Instead, our goal is to point out the aspects of the modified SPA that it shares with the pre-trade transparent trading regime from Section 3, as well as those that allow it to outperform that regime, hoping to identify economic principles that likely hold beyond our simple model.

First, by alleviating the winner's curse for the uninformed type, the SPA pricing rule allows uninformed dealers to quote more aggressively than in the first-price auction. This in turn makes it less attractive for the low type to pretend to be uninformed, and hence allows higher probabilities of trade to be supported when both dealers are uninformed. However, the SPA pricing rule makes quotes less attractive for the customer in case one of the dealers is uninformed, which may prevent efficient trade. This issue is resolved by adding the post-auction bargaining stage that allows the winning dealer and the customer to reach mutually agreeable terms of trade. Thus, our results provide a theoretical support for allowing post-auction bargaining on trading platforms in settings with adverse selection.

The introduction of a fee paid by the customer allows us to implement an interior probability of trade conditional on both dealers being uninformed when $\Delta < 1/2$. (The fee is not needed when $\Delta \geq 1/2$.) Intuitively, the fee is high enough that the customer may sometimes decline trading with the uninformed type, which then reduces the low type's incentives to pretend to be uninformed. While this intuition is more difficult to apply to real-life trading mechanisms, one conclusion for policy is that the presence of trading fees in adverse selection environments may not necessarily be a sign of inefficiency—eliminating or lowering these fees may lead to worse outcomes due to subtle informational frictions.

Another feature of the trading protocol from Theorem 5 is that the optimal mechanism offers transparency "within" each trade, but no transparency "across trades." This echoes the intuition we emphasized when discussing pre-trade versus full transparency: Once dealers' information is symmetrized (through seeing each other's quotes), there is no point in trying to make the customer symmetrically informed as well; in fact, this attempt can backfire as it gives the dealers an incentive to manipulate market beliefs. We discuss broader economic implications of that observation in Section 5.

Unlike in Section 3, the insights derived in this section do not require the additional assumption on the dealers' patience and asset value persistence given in (2.1). In particular, the optimal mechanism can always sustain a fully separating equilibrium, which may be a desirable property from the perspective of the information-aggregation role of prices. However, a caveat of the modified SPA is that it may have some less efficient equilibria. In Appendix OA.4, we construct an alternative trading regime that achieves unique implementation (under the weaker solution concept from Section 3) of the optimal allocation. That trading regime is also a variant of the SPA with further modifications of the payments, a coarse price grid, and rationing by the platform. These modifications render the design less practical, highlighting a tension between simplicity of trading rules and robustness to equilibrium selection.

5 Concluding remarks

Our paper provides a framework for studying several prominent trading regimes, as well as the optimal trading protocol, in a model with two-sided adverse selection inspired by real-life OTC markets. The analysis leads to a number of policy implications, such as the superiority of pre-trade transparent markets over post-trade and fully transparent markets and the potential efficiency-enhancing role of second-price auctions and post-auction bargaining. In order to tractably analyze a variety of trading protocols, we made a number of stylized assumptions about the setting, which could raise concerns about the credibility and robustness of these insights. In this last section, we briefly discuss some extensions as well as limitations of our model.

Desirability of post-trade transparency. One of the main insights of our paper is that post-trade transparency has limited value in ensuring market efficiency, assuming the designer has access to other forms of transparency. Post-trade transparency is welfare-dominated by pre-trade transparency, and can only reduce welfare when added to a pre-trade transparent market. The optimal mechanism we constructed in Section 4 also does not feature any post-trade transparency.

The mechanism-design perspective sheds light on the fundamental reasons for the ineffectiveness of post-trade transparency. As we have explained in Section 4.1, if the designer has full control over the market protocol, then post-trade transparency is at best redundant as a consequence of the general revelation principle for dynamic interactions (see Myerson, 1982, and the related idea of the *inscrutability principle* in Myerson, 1983). In the language of indirect trading mechanisms, disclosure of endogenous variables—such as the price—creates incentives for agents to manipulate them for their private benefit. Consequently, as long as the designer can choose the trading protocol within each period, disclosing information about past trades only makes the design problem more difficult.[36]

[36] Another consequence of the same basic principle of pooling incentives is that forcing dealers to provide

That being said, post-trade transparency may be desirable when the designer has limited control over the market protocol. For example, implementing pre-trade transparency may be politically costly in some markets due to the concentrated interest of the dealers. (Indeed, dealer profits generally decrease when pre-trade transparency is introduced in our framework). Regulators may also be hesitant to discontinuously change the trading regime due to the risk of the regulation's unintended consequences in markets where the trading protocol evolved organically over years to balance heterogeneous needs of various market participants.[37] In such cases, by Theorem 1, introducing post-trade transparency may be a more "cautious" way of increasing efficiency relative to a fully opaque market.

Post-trade transparency may have benefits not captured by our model. Duffie et al. (2017) show that post-trade transparency (introduction of a benchmark) allows traders to make more efficient entry decisions (primarily by avoiding costly entry when gains from trade are small).[38] Additionally, post-trade transparency may be an important source of information about the market for actors not directly interested in trading, such as regulators or real-economy decision makers. While this last benefit could be substantial, we note that our critique of post-trade transparency applies only to *immediate* public disclosure of trade data (the belief manipulation effect is present only when dealers' information is relevant for transactions following the disclosure). Our results are fully consistent with the idea that it can be beneficial to disclose trade data with a (small) time lag—and indeed, our focus on separating equilibria means that such "lagged" post-trade transparency would have the same social value in all the regimes we have studied.

Finally, another potentially important margin is traders' endogenous choice among trading venues. Recently, Wu (2025) developed a dynamic model in which traders choose between a transparent centralized market and an opaque decentralized market. In her framework, transparency affects welfare not only through trading outcomes within a venue, but also by altering subsequent venue choices and thereby the evolution of publicly available price information. Incorporating this margin into our dealer–customer framework would be an interesting direction for future research.

Multiple dealers. Throughout the paper, we restricted attention to two dealers. OTC markets tend to exhibit a core-periphery structure (see, for example, Wang, 2016, and the

two-sided quotes may enhance welfare, as long as dealers remain uninformed about the customer's desired direction of trade when committing to quotes (an assumption that may be violated in many OTC markets that lack trader anonymity, see, e.g., Duffie, 2012 and Lee and Wang, 2024). In our framework, introduction of two-sided quotes would lead to existence of an efficient equilibrium but it would generally not increase welfare in the worst equilibrium—we leave a comprehensive analysis of two-sided quotes for future research.

[37] This was a concern for the UK regulators when formulating policy recommendations, see FCA (2024).

[38] Because pre-trade transparency shifts more rents towards customers, it is not clear which form of transparency would lead to more efficient entry in a model with costly entry and stochastic gains from trade.

references therein), so assuming that a small number of dealers intermediate most of the trades with customers is realistic. Adding more dealers to the model would complicate the notation without changing our insights qualitatively. In the opaque and post-trade transparent markets, equilibria would not be affected. In the pre-trade transparent market, equilibria would likely become more efficient with large N but adverse selection for uninformed dealers would keep welfare below the efficient level.

Interestingly, even when many dealers are available to provide quotes on OTC trading platforms, auctions tend to involve a small number of dealers. For example, Riggs et al. (2020) document that only about three dealers are contacted in a typical trade on the two largest dealer-to-customer SEFs for credit default swaps. Wang (2023) constructs a theoretical model in which dealers can strategically choose to ignore a request for quote, and shows that, in equilibrium, the customer contacts only two dealers.

Persistence in the asset value. To simplify our analysis, we assumed that the asset value is persistent over at most two periods. We could instead allow it to be persistent over a longer horizon, e.g., by assuming that new information arrives to the market according to a more general stochastic process. We believe that such a model would be less tractable without offering qualitatively new insights. In post-trade transparent regimes, a single informative transaction would lead to efficient trading until new information arrives to the market. Thus, it may seem that this version of our model would make post-trade transparency more attractive. That need not be the case, however, because more persistence would also exacerbate dealers' incentives to manipulate market beliefs (a one-time price manipulation would have a more long-lasting effect on beliefs of future customers). Even if the post-trade transparent market became more efficient with more persistence in the asset value, our intuition for Theorem 2 suggests that it would still be dominated by pre-trade transparency: As long as the auction symmetrizes dealers' information after the first trade following the arrival of news to the market, the subsequent trading rounds would feature fully efficient trade. Finally, our optimal dynamic mechanism was designed to achieve high trade volume within each period and did not rely on intertemporal incentives, so we would expect a similar mechanism to emerge as optimal in environments with higher asset value persistence.

Information acquisition. Throughout the paper, we treated the information structure as exogenous. It would be interesting to endogenize the information structure by letting dealers choose the probability q of being informed optimally at an ex-ante stage subject to a cost. It is well known that more competitive market protocols—such as our pre-trade transparent market—may have adverse effects on the incentives of market participants to acquire information (the insight of Grossman and Stiglitz (1980) is perhaps the most famous example of this phenomenon). More recently, Glode and Opp (2020) and Baldauf and

Mollner (2020) showed that differential incentives to acquire information may affect welfare comparisons of trading protocols. Since inefficiencies in our model vanish at both extremes—with fully informed as well as fully uninformed dealers—the effects of information acquisition on welfare are not ex-ante obvious. We leave this direction for future research.

Semi-pooling equilibria. A limitation of our analysis was the assumption that dealers' quotes reveal their information on equilibrium path—this assumption rules out semi-pooling equilibria which could potentially be more efficient than separating equilibria. Characterizing all semi-pooling equilibria appears to be a daunting task.[39] If we allowed $\Delta \geq 1$, then all regimes would have a fully efficient pooling equilibrium and thus the *same* maximal welfare. This suggests that allowing for some pooling when $\Delta < 1$ makes different trading regimes more *similar* without necessarily changing the welfare ranking.

Privately informed customers. Our model—with dealers having an informational advantage over customers—corresponds to markets where customers trade primarily for liquidity reasons. In practice, customers may have information about the asset value and trade to make a profit against an uninformed dealer. Consider a simple extension of our framework in which a fraction of customers have no liquidity need but know the true value of the asset, and dealers can infer the customer's trading motive by observing her identity (similarly to the model of Lee and Wang, 2024). In the informed-customer segment of the market, in all regimes we studied, the equilibrium outcome would be efficient except in the case when the asset value is low and the customer only observes quotes from uninformed dealers. The introduction of post-trade transparency would be beneficial because it would help mitigate the adverse selection problem faced by the uninformed dealer. Pre-trade transparency would additionally increase the likelihood that the customer receives at least one quote from an informed dealer, so it would weakly dominate post-trade transparency. Thus, in the simple extension, our results would continue to hold. A more realistic model in which dealers cannot perfectly infer the trading motive of the customer would be difficult to analyze.

Inventory management and the inter-dealer market. Our model focuses on the dealer-to-customer segment of OTC markets. However, a large fraction of trades happen between dealers trying to rebalance their positions—an aspect that our framework does not capture due to the simplifying assumption of unit trade and constant marginal values. Before the introduction of TRACE, dealers expressed a concern that too much transparency would hurt their ability to intermediate trades by raising their price impact in the inter-dealer market (Asquith et al., 2019). And indeed, dealers might prefer to use non-transparent

[39] See Dworczak (2015) for an attempt to characterize a subset of semi-pooling equilibria in a model featuring similar frictions.

mechanisms such as dark pools and size-discovery sessions to avoid price impact (see, for example, Zhu, 2012 and Duffie and Zhu, 2017). Of course, dealers' opposition to transparent mechanisms does not imply that transparency is not welfare-enhancing overall (for example, dealers in our model would lose from a transition from the opaque to the pre-trade transparent market). The question is whether dealers' reluctance to take on large positions in the face of a transparent inter-dealer market could reduce market efficiency.

Answering this question would require a model with non-linear dealer values and an inter-dealer market, such as the one proposed by Babus and Parlatore (2022)—an extension of the linear-quadratic double-auction model of Rostek and Weretka (2012). Unfortunately, the tractability of the double-auction model (and—to the best of our knowledge—most financial market models featuring a meaningful notion of trade size) is limited to linear equilibria and Gaussian information structures that are typically not preserved under interesting modifications of market transparency.[40] Thus, understanding optimal market transparency in the presence of inter-dealer markets remains an important open problem.

References

ANTILL, S. AND D. DUFFIE (2020): "Augmenting Markets with Mechanisms," The Review of Economic Studies, 88, 1665–1719.

AQUILINA, M., G. IBIKUNLE, V. MOLLICA, AND T. STEFFEN (2022): "The visible hand: benchmarks, regulation, and liquidity," Journal of Financial Markets, 61, 100734.

ASQUITH, P., T. COVERT, AND P. A. PATHAK (2019): "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market," Available at SSRN 2320623.

ASRIYAN, V., W. FUCHS, AND B. GREEN (2017): "Information Spillovers in Asset Markets with Correlated Values," American Economic Review, 107, 2007–40.

BABUS, A. AND C. PARLATORE (2022): "Strategic fragmented markets," Journal of Financial Economics, 145, 876–908.

BACK, K., R. LIU, AND A. TEGUIA (2020): "Signaling in OTC Markets: Benefits and Costs of Transparency," Journal of Financial and Quantitative Analysis, 55, 47–75.

BAGWELL, K. AND G. RAMEY (1991): "Oligopoly limit pricing," The Rand Journal of Economics, 155–172.

BALDAUF, M. AND J. MOLLNER (2020): "High-Frequency Trading and Market Performance," The Journal of Finance, 75, 1495–1526.

BANKS, J. S. AND J. SOBEL (1987): "Equilibrium selection in signaling games," Econometrica: Journal of the Econometric Society, 647–661.

[40]The only exception we are aware of is Ollar et al. (2021) but even that paper can only analyze a special class of transparency regimes.

Baron, D. P. and D. Besanko (1984): "Regulation and information in a continuing relationship," Information Economics and policy, 1, 267–302.

Bergemann, D. and J. Hörner (2018): "Should first-price auctions be transparent?" American Economic Journal: Microeconomics, 10, 177–218.

Bessembinder, H. and W. Maxwell (2008): "Markets: Transparency and the corporate bond market," Journal of Economic Perspectives, 22, 217–234.

Blouin, M. R. and R. Serrano (2001): "A decentralized market with common values uncertainty: Non-steady states," The Review of Economic Studies, 68, 323–346.

Budish, E., P. Cramton, and J. Shim (2015): " The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response *," The Quarterly Journal of Economics, 130, 1547–1621.

Chen, D. (2024): "Optimal Exchange Design," Working Paper.

Cho, I.-K. and D. M. Kreps (1987): "Signaling games and stable equilibria," The Quarterly Journal of Economics, 102, 179–221.

Crémer, J. and R. P. McLean (1985): "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent," Econometrica, 53, 345–362.

Daley, B. and B. Green (2012): "Waiting for News in the Market for Lemons," Econometrica, 80, 1433–1504.

Diamond, P. A. (1971): "A model of price adjustment," Journal of Economic Theory, 3, 156–168.

Duffie, D. (2012): Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets, Princeton University Press.

Duffie, D., P. Dworczak, and H. Zhu (2017): "Benchmarks in Search Markets," The Journal of Finance, 72, 1983–2044.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-Counter Markets," Econometrica, 73, 1815–1847.

Duffie, D. and J. C. Stein (2015): "Reforming LIBOR and Other Financial Market Benchmarks," Journal of Economic Perspectives, 29, 191–212.

Duffie, D. and H. Zhu (2017): "Size Discovery," The Review of Financial Studies, 30, 1095–1150.

Dworczak, P. (2015): "The Effects of Post-Auction Bargaining between Bidders," Working Paper, Stanford University.

——— (2020): "Mechanism Design With Aftermarkets: Cutoff Mechanisms," Econometrica, 88, 2629–2661.

Dye, R. A. (1985): "Disclosure of Nonproprietary Information," Journal of Accounting Research, 23, 123–145.

Edwards, A. K., L. E. Harris, and M. S. Piwowar (2007): "Corporate bond market transaction costs and transparency," The Journal of Finance, 62, 1421–1451.

Fama, E. F. and M. H. Miller (1972): The Theory of Finance, Holt, Rinehart, and Winston, New York.

FCA (2024): "Policy Statement for Improving transparency for bond and derivatives markets," Report.

Fuchs, W., A. Öry, and A. Skrzypacz (2016): "Transparency and distressed sales under asymmetric information," Theoretical Economics, 11, 1103–1144.

Glebkin, S., B. Z. Yueshen, and J. Shen (2023): "Simultaneous multilateral search," The Review of Financial Studies, 36, 571–614.

Glode, V. and C. C. Opp (2020): "Over-the-counter versus limit-order markets: The role of traders' expertise," The Review of Financial Studies, 33, 866–915.

Glosten, L. R. and P. R. Milgrom (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," Journal of Financial Economics, 14, 71–100.

Grossman, S. J. and J. E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," The American Economic Review, 70, 393–408.

Guerrieri, V. and R. Shimer (2014): "Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality," American Economic Review, 104, 1875–1908.

Guerrieri, V., R. Shimer, and R. Wright (2010): "Adverse selection in competitive search equilibrium," Econometrica, 78, 1823–1862.

Gul, F. and W. Pesendorfer (2012): "The war of information," The Review of Economic Studies, 79, 707–734.

Halac, M. and I. Kremer (2020): "Experimenting with career concerns," American Economic Journal: Microeconomics, 12, 260–88.

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

Hörner, J. and N. Vieille (2009): "Public vs. private offers in the market for lemons," Econometrica, 77, 29–69.

Hou, D. and D. Skeie (2013): "LIBOR: Origins, Economics, Crisis, Scandal and Reform," in The New Palgrave Dictionary of Economics, Online Edition, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan.

Janssen, M., P. Pichler, and S. Weidenholzer (2011): "Oligopolistic Markets with Sequential Search and Production Cost Uncertainty," RAND Journal of Economics, 42, 444–470.

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

——— (2022): "Public vs. Private offers with informed and forward-looking dealers," Working Paper.

Kargar, M., B. R. Lester, S. Plante, and P.-O. Weill (2024): "Sequential Search for Corporate Bonds," Available at SSRN: https://ssrn.com/abstract=4284242.

Kaya, A. and K. Kim (2018): "Trading dynamics with private buyer signals in the market for lemons," The Review of Economic Studies, 85, 2318–2352.

Kaya, A. and S. Roy (2022a): "Market screening with limited records," Games and economic behavior, 132, 106–132.

——— (2022b): "Price Transparency and Market Screening," Working paper.

——— (2024): "Repeated Trading: Transparency and Market Structure," American Economic Review, 114, 2388–2435.

Kim, K. (2017): "Information about sellers' past behavior in the market for lemons," Journal of Economic Theory, 169, 365–399.

Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," Econometrica, 53, 1315–1335.

Lauermann, S. and A. Wolinsky (2016): "Search with adverse selection," Econometrica, 84, 243–315.

Lee, T. and C. Wang (2024): "Regulating Over-the-Counter Markets," Journal of Finance (forthcoming).

Lester, B., A. Shourideh, V. Venkateswaran, and A. Zetlin-Jones (2019): "Screening and adverse selection in frictional markets," Journal of Political Economy, 127, 338–377.

Loertscher, S., E. V. Muir, and P. G. Taylor (2022): "Optimal market thickness," Journal of Economic Theory, 200, 105383.

Meli, J., Z. Todorova, and A. Diaz (2024): "Trade Protocols, Transparency, and Liquidity–Theory and Evidence from the European Corporate Bond Market," Available at SSRN.

Milgrom, P. and J. Mollner (2021): "Extended proper equilibrium," Journal of Economic Theory, 194, 105258.

Myerson, R. B. (1982): "Optimal coordination mechanisms in generalized principal-agent problems," Journal of Mathematical Economics, 10, 67–81.

——— (1983): "Mechanism Design by an Informed Principal," Econometrica, 51, 1767–1797.

Ollar, M., M. J. Rostek, and J. H. Yoon (2021): "Privacy in Markets," Available at SSRN: https://ssrn.com/abstract=3071374.

Ordin, A. and R. Sverchkov (2024): "Incentives to Lose: Disclosure of Cover Bids in OTC Markets," Available at SSRN.

Osborne, M. J. and A. Rubinstein (1990): Bargaining and markets, Academic press.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu (2020): "Swap trading after Dodd-Frank: Evidence from index CDS," Journal of Financial Economics, 137, 857–886.

Rostek, M. and M. Weretka (2012): "Price Inference in Small Markets," Econometrica, 80, 687–711.

Rostek, M. and X. Wu (2021): "Improving Access to Information Through Market Design," Working paper, available at SSRN 3974087.

Schmidt, K. M. (1993): "Commitment through incomplete information in a simple repeated bargaining game," Journal of Economic Theory, 60, 114–139.

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

Vairo, M. and P. Dworczak (2024): "Optimal Transparency in Financial Over-the-Counter Markets," Working Paper.

Vickrey, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," The Journal of Finance, 16, 8–37.

Vida, P. and T. Honryo (2021): "Strategic stability of equilibria in multi-sender signaling games," Games and Economic Behavior, 127, 102–112.

Wang, C. (2016): "Core-Periphery Trading Networks," Available at SSRN: https://ssrn.com/abstract=2747117.

——— (2023): "The limits of multi-dealer platforms," Journal of Financial Economics, 149, 434–450.

Wolinsky, A. (1990): "Information revelation in a market with pairwise meetings," Econometrica: Journal of the Econometric Society, 1–23.

Wu, X. (2025): "Dynamic Market Choice," Working paper.

Zhang, A. L. (2022): "Competition and manipulation in derivative contract markets," Journal of Financial Economics, 144, 396–413.

Zhu, H. (2012): "Finding a Good Price in Opaque Over-the-Counter Markets," The Review of Financial Studies, 25, 1255–1285.

A Proofs

This section collects the proofs of all the main results of the paper.

Throughout our proofs—unless explicitly stated otherwise—we focus on the case when the customer is a buyer. The case when the customer is a seller is symmetric (with the low-type dealer playing the role of the high-type dealer and *vice versa*, and prices flipping their signs). Explicit treatment of the customer-seller case is only needed in two scenarios. First, if some information about the value of the asset was revealed in period $t-1$ and $Z_t = 1$, then continuation play depends on whether the customer arriving in period t trades in the same direction as the customer who arrived in period $t-1$. Second, one must take both cases into account to correctly compute posterior beliefs (e.g., when no transaction took place in the previous period and the state is persistent).

All proofs rely on the fact—established formally in the next subsection—that, under our assumptions, the set of stationary equilibrium strategies in the trading regimes that we study coincides with those of a simpler two-period game.

A.1 Two-period interpretation of the model

Consider the following two-period game. In period zero, v is drawn, dealers privately observe their signals, and the trading game is played. With probability $1-\gamma$, the game ends after period 0. With probability γ, the game proceeds to period one, in which dealers play the trading game again against a new customer, with the value of v remaining fixed and dealers observing no further signal realizations. Dealers discount the future by a factor δ.

The set of equilibrium strategies of the two-period game and our original infinite-horizon games are equivalent. The equivalence holds because, under the stationarity restriction, dealers' strategies only depend on their private information and the realization of $Z_t \in \{0, 1\}$. Since all variables are i.i.d. redrawn when $Z_t = 0$, the game "resets" in this event, in the sense that the resulting continuation game is history-independent. Moreover, in that game, dealers only contemplate their current payoffs and their continuation payoffs conditional on the state remaining fully persistent in the following period, and they will effectively "discount" this continuation payoff by a factor of $\delta\gamma$ (the discount factor times the probability of reaching this stage of the game). When $Z_t = 1$, players know with certainty that the game will reset in the following period, and therefore behave as if this was the final period of the game. As a result, stationary equilibrium strategies in the infinite-horizon game when $Z_t = 0$ and $Z_t = 1$ are the same as equilibrium strategies in the two-period game in period zero and period one, respectively.

In addition, the following lemma shows that, regardless of the trading regime, discounted expected equilibrium welfare in the infinite-horizon game is proportional to its counterpart in the two-period game. Thus, the welfare rankings across trading regimes are the same in the two games.

Lemma 1. *Equilibrium welfare under protocol r is*

$$\Lambda_r = (1-\delta)\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\Lambda_r(Z_t)] = \frac{\Lambda_r(0) + \delta\gamma\Lambda_r(1)}{1+\delta\gamma},$$

where $\Lambda_r(Z)$, $Z \in \{0,1\}$, are equilibrium probabilities of trade conditional on $Z_t = Z$.

Proof. Let $p_t = (\Pr(Z_t = 0), \Pr(Z_t = 1))$, be the ex ante distribution of Z_t. Since Z_t follows a Markov chain, we can write

$$p_t = p_0 \times P^t,$$

where $p_0 = (1,0)$ is the distribution of Z_0 and

$$P \equiv \begin{pmatrix} 1-\gamma & \gamma \\ 1 & 0 \end{pmatrix}$$

is the Markov transition matrix of Z_t. P can be decomposed as

$$P = A \begin{pmatrix} 1 & 0 \\ 0 & -\gamma \end{pmatrix} A^{-1}, \quad A = \begin{pmatrix} 1 & -\gamma \\ 1 & 1 \end{pmatrix}.$$

Hence,

$$p_t = p_0 \times A \begin{pmatrix} 1 & 0 \\ 0 & (-\gamma)^t \end{pmatrix} A^{-1} = \left(\frac{1+\gamma(-\gamma)^t}{1+\gamma}, \frac{\gamma(1-(-\gamma)^t)}{1+\gamma} \right).$$

Given $\Lambda_r(Z)$, the discounted expectation can be written as

$$\sum_{t=0}^{\infty} \delta^t p_t \times \begin{pmatrix} \Lambda_r(0) \\ \Lambda_r(1) \end{pmatrix} = \frac{\Lambda_r(0) + \delta\gamma\Lambda_r(1)}{(1-\delta)(1+\delta\gamma)}.$$

□

Given the weights that arise from Lemma 1, we define

$$\alpha \equiv \frac{1}{1+\delta\gamma}. \tag{A.1}$$

With this equivalence in mind, throughout the Appendix, we refer to realizations of $Z_t = z \in \{0,1\}$ simply as period z. In particular we will often refer to "period zero" and "period one" omitting the physical time index t.

A.2 Equilibrium existence

Throughout, we restrict attention to parameter values satisfying Assumption (2.1) under which separating equilibria exist in *all* of the regimes that we study.

Lemma 2. *Let* $\overline{\delta}_{\text{exist}} \equiv \min\{\overline{\delta}_1, \overline{\delta}_2\}$*, where*

$$\overline{\delta}_1 \equiv \frac{2[(2-q)\Delta^2 - q]}{(1-q)[(1+\Delta)\max\{1, 2\Delta\} + 2\Delta]},$$
$$\overline{\delta}_2 \equiv \frac{4(1-\Delta^2)(\Delta+q)}{(1+\Delta)(1-q)[(1+\Delta)\max\{1, 2\Delta\} + 2(1-\Delta)]}.$$

A separating equilibrium (satisfying all refinements imposed in Section 2) exists in all of the regimes that we study if $q < 2\Delta^2/(1+\Delta^2)$ *and* $\delta\gamma \leq \overline{\delta}_{\text{exist}}$*.*

One can verify that the assumption $q < 2\Delta^2/(1+\Delta^2)$ is equivalent to $\overline{\delta}_{\text{exist}} > 0$. The lemma will be proven separately for each regime.

A.3 Proof of Theorem 1

Throughout, we rely on results proven in Appendix A.1. We first characterize equilibria in the opaque market, then characterize equilibria in the post-trade transparent market, and finally compare the equilibrium welfare levels across these two regimes.

A.3.1 The opaque market

In Proposition 1, we characterize equilibrium outcomes of the opaque market. The set of equilibrium outcomes in this framework coincide with those in the static monopoly benchmark described in Section 3.1.

Proposition 1 (Structure of equilibria in the opaque market)**.** *In all equilibria, in each period* t*, equilibrium outcomes in the opaque regime are payoff-equivalent to the static monopoly trading pattern: The low type quotes* $-1+\Delta$ *which is accepted with probability one, the uninformed type quotes* Δ *which is accepted with probability* $\lambda_u \in [0, \bar{\lambda}_u]$ *for some* $\bar{\lambda}_u < 1$*, and the high type quotes* $1+\Delta$ *which is accepted with some probability* $\lambda_h \in [0, \bar{\lambda}_h]$ *for some* $\bar{\lambda}_h < \bar{\lambda}_u$*.*

In the least efficient equilibrium, the uninformed and the high-type dealers do not trade. In the most efficient equilibrium, in each period, the uninformed type trades with probability $\Delta/(1+\Delta)$*, and the high type trades with probability* $(\Delta/(1+\Delta))^2$*.*

Proof. In an opaque regime, dealers do not learn anything about each other's type, which makes the game starting in every period a repetition of the static one. We begin by characterizing dealers' quoting strategies. Let b_σ be any price in the support of the equilibrium quoting strategy of a dealer with type $\sigma \in \{l, u, h\}$, and let $\underline{b}_\sigma, \bar{b}_\sigma$ be the respective infimum and supremum of the support.

First, note that $\bar{b}_l \leq -1+\Delta$; otherwise, a quote equal to $\bar{b}_l$ is accepted by the customer with zero probability, and the low-type dealer could do strictly better by offering $-1+\Delta-\epsilon$ (which is always accepted by the customer provided that $\epsilon > 0$ is small enough). Second, we show that $\underline{b}_l \geq -1+\Delta$. Towards a contradiction, suppose that $\underline{b}_l < -1+\Delta$. If this is the case, the low-type dealer can profitably deviate to $\tilde{b} \in (\underline{b}_l, -1+\Delta)$: The customer has a strict incentive to accept this quote and the dealer gets to sell at a higher price, a contradiction. Summarizing, we have $\underline{b}_l = \bar{b}_l = -1+\Delta$. Additionally, a quote of $-1+\Delta$ must be accepted with probability one in equilibrium. If this were not the case, then the low type could profitably deviate to $-1+\Delta-\epsilon$ which the customer would accept with probability one, leading to a strict improvement in the dealer's payoff for ϵ small enough, a contradiction.

By the argument provided in the main text, it is without loss of generality to assume that the uninformed dealer quotes $b_u = \Delta$. In particular, if the uninformed type quoted $0 \leq b_u < \Delta$, the customer would agree to trade for sure; but this would give rise to a profitable deviation for the low type—by quoting b_u, she could trade for sure (conditional on being contacted) at a strictly higher price than in equilibrium. A similar argument establishes that $b_h = 1+\Delta$.

Next, we argue that satisfying the incentive constraints ($\mathrm{IC}_{l\to u}$), ($\mathrm{IC}_{u\to h}$) is sufficient for equilibrium. We can specify off-path beliefs for the customer by making them as low as possible given the belief-monotonicity refinement: For any quote b, the customer believes that the expected asset value is 1 (quoting dealer has a high type) when $b \geq 1+\Delta$, 0 (quoting dealer is uninformed) when $b \in [\Delta, 1+\Delta)$, and -1 (quoting dealer has a low type) if $b < \Delta$. For future reference, we refer to this system of beliefs as *the most pessimistic opaque-market beliefs.* Then, it is easy to check that no dealer type has a profitable deviation, and that each customer is making optimal decisions given her belief. In particular, this establishes equilibrium existence for this regime (see Lemma 2).

Finally, letting $\lambda_u(Z)$ and $\lambda_h(Z)$ denote the equilibrium probabilities that the uninformed and the high-type dealers trade conditional on $Z_t = Z$, it follows from ($\mathrm{IC}_{l\to u}$) and ($\mathrm{IC}_{u\to h}$) that $\lambda_h(Z) = \lambda_u(Z) = 0$ for all $Z \in \{0,1\}$ in the least efficient equilibrium, and $\lambda_u(Z) = \frac{\Delta}{1+\Delta}$ and $\lambda_h(Z) = \left(\frac{\Delta}{1+\Delta}\right)^2$ for all $Z \in \{0,1\}$ in the most efficient one. $\square$

A.3.2 The post-trade transparent market

Proposition 2 (Structure of equilibria in the post-trade transparent market)**.** *In all equilibria of the post-trade transparent market, if $Z_t = 0$, the low type quotes $-1+\Delta$ which is accepted with probability one, the uninformed type quotes Δ which is accepted with some probability $\lambda_u \in (0,1)$, and the high type quotes $1+\Delta$ which is accepted with some probability*

$\lambda_h \in (0,1)$. *If* $Z_t = 1$*:*

- *If a low-type or high-type dealer transacted in period* $t-1$*, the equilibrium is fully efficient in period* t*;*
- *If there was no transaction or if an uninformed dealer transacted in period* $t-1$*, then equilibrium outcomes conform to the static monopoly trading pattern described in Proposition 1.*

Proof. We proceed backwards by first deriving continuation equilibria in period one, and then using this result to characterize equilibrium strategies in period zero.
Step 1: Continuation equilibrium in period one. There are three kinds of on-path period-one histories to consider determined by the transaction or lack thereof announced at the end of period zero. First, if an informed dealer (either with high or low type) transacted in period zero, then, because equilibrium quotes are fully revealing, the market becomes symmetrically informed about the value of the asset. Under the never-dissuaded-once-convinced refinement, the game that ensues is one of complete information in which the unique equilibrium features trade with probability one at the monopolistic price. In this equilibrium, conditional on being contacted, both dealers provide the same quote, and therefore the customer is indifferent between trading with either one of them. Under our symmetry and stationarity assumptions, the period-one customer randomizes which dealer to contact first with probability 1/2.

Second, let us look at the continuation game after an uninformed-type dealer, say dealer A, transacted in period zero. Dealer A learns nothing in period zero about the value of the asset. Dealer B learns that her competitor is uninformed with probability one, but her estimate of the value of the asset remains unchanged. Because trade is anonymous, the period-one customer enters the market with symmetric beliefs about dealers' types with full support on $\{l, u, h\}$. By the same argument as in Proposition 1, conditional on being contacted, a dealer with belief $v' \in \{-1, 0, 1\}$ is going to quote $v' + \Delta$, and the low-type dealer B trades with probability one. To ensure that the uninformed type does not want to mimic the high type and that the low type does not want to mimic the uninformed type, we need precisely the same conditions as we derived for the opaque market. That is, letting $\lambda_b^u(b_u), \lambda_b^h(b_u)$ be, respectively, the probabilities that the quote of the uninformed and the high types is accepted by the customer-buyer at this node of the game, we must have

$$\lambda_b^u(b_u) \leq \frac{\Delta}{1+\Delta} \quad \text{and} \quad \lambda_b^h(b_u) \leq \lambda_b^u(b_u)\frac{\Delta}{1+\Delta}.$$

A symmetric condition must be satisfied for the customer-seller's acceptance probabilities conditional on an uninformed and low type, which we denote by $\lambda_s^u(b_u)$ and $\lambda_s^l(b_u)$. These conditions are also sufficient to rule out profitable deviations for both dealers if we set off-path beliefs equal to the most pessimistic opaque-market beliefs, as defined in the proof of Proposition 1.

Third, let us look at the continuation equilibrium following no trade in period zero. Because of symmetry (across asset values and desired positions of the customer), this event is equally likely to happen when the asset value is either high or low, so it is uninformative about the value of the asset. Hence, dealers enter this part of the game with the same belief about the asset value that they had in period zero, and the period-one customer's belief about dealers' types continues to be supported on $\{l, u, h\}$. The game is then the same as in the static opaque market, and the set of equilibrium outcomes are described by the static monopoly trading pattern. In particular, if the customer is a buyer, the uninformed and the high-type dealers trade with probabilities $\lambda_b^u(nt), \lambda_b^h(nt)$, satisfying $\lambda_b^u(nt) \leq \frac{\Delta}{1+\Delta}$ and $\lambda_b^h(nt) \leq \frac{\Delta}{1+\Delta}\lambda_b^u(nt)$. The relevant trading probabilities in the customer-seller case, $\lambda_s^u(nt)$ and $\lambda_s^l(nt)$, are defined symmetrically.

Step 2: Equilibrium strategies in period zero. For $\sigma \in \{l, u, h\}$, let $\underline{b}_\sigma$ and $\overline{b}_\sigma$ be the infimum and the supremum of the period-zero quoting strategy of a dealer with type σ. First, the low type must play a pure strategy on $b_l = -1 + \Delta$, and this quote must be accepted with probability one. To see this, first observe that it must be that $\overline{b}_l \leq -1 + \Delta$, for if not, the low type strictly benefits from deviating to $-1 + \Delta - \epsilon$ where $\epsilon > 0$ is sufficiently small. This deviation leads to a total payoff of at least Δ, whereas the payoff from quoting $\overline{b}_l$ is equal to the continuation payoff conditional on no trade, which is bounded above by $\delta\gamma\Delta/2 < \Delta$. If $\underline{b}_l < \overline{b}_l$, then the low type strictly benefits from deviating to $b \in (\underline{b}_l, \overline{b}_l)$: This quote is strictly better than $\underline{b}_l$ as it leads to trade with probability one in period zero and does not affect continuation payoffs (by belief monotonicity). Thus, $\underline{b}_l = \overline{b}_l \equiv b_l \leq -1 + \Delta$. Moreover, if $b_l < -1 + \Delta$, then the low type would benefit from deviating to a quote slightly below $-1 + \Delta$, which still leads to trade with probability one in the current period, without affecting continuation payoffs (due to stationarity and part (ii) of the monotonicity refinement). Hence, $b_l = -1 + \Delta$. Finally, if the customer accepts $-1 + \Delta$ with probability less than one, then the low type can strictly benefit from deviating to $-1 + \Delta - \epsilon$ with $\epsilon > 0$ sufficiently small, which the customer would strictly want to accept and which does not affect continuation payoffs due to belief monotonicity. Thus, the customer must accept the low type's quote with probability one.

Second, the high type must quote $\underline{b}_h = \overline{b}_h = 1 + \Delta$. Suppose that the high type provides a quote $b_h < 1 + \Delta$ in equilibrium. This quote must be accepted with probability one by the customer and lead to a continuation payoff of $\Delta/2$. Consider now the high type's payoff from deviating to $\tilde{b} = b_h + \epsilon < 1 + \Delta$. By monotonicity of beliefs and stationarity of equilibrium strategies, this quote is accepted with probability one by the current customer and leads to the same continuation payoff as b_h (since it leads to the same hierarchy of beliefs), a contradiction. Hence, $b_h \geq 1 + \Delta$. Moreover, any equilibrium with $b_h > 1 + \Delta$ is outcome-equivalent to one where $b_h = 1+\Delta$ and the customer rejects with probability one, so without loss of generality we can set $b_h = 1 + \Delta$.

Third, we argue that, if the uninformed dealer quoted $b_u < \Delta$, then either the uninformed or the low type would have a profitable deviation. To see this, note that the low-type dealer's equilibrium payoff (conditional on being contacted in period zero) is $\Delta(1 + \delta\gamma/2)$. If the low type deviates to b_u, she gets to trade with probability one in the current period and ensures a continuation payoff of at least $\delta\gamma\Delta/4$. Hence, equilibrium requires that

$$\Delta\left(1 + \frac{\delta\gamma}{2}\right) \geq b_u + 1 + \frac{\delta\gamma}{4}\Delta \iff b_u \leq \Delta - 1 + \frac{\delta\gamma}{4}\Delta.$$

Now, consider the uninformed dealer's payoff from deviating to $\tilde{b}_u = \Delta - \epsilon > b_u$, where $\epsilon > 0$ is sufficiently small. By monotonicity of beliefs, this quote is accepted with certainty by the current customer and leads to a non-negative continuation payoff. In equilibrium, on the other hand, the uninformed dealer is trading with probability one in the current period at a price equal to b_u, and receives a continuation payoff bounded above by $\delta\gamma\Delta/2$. Hence, a necessary condition to rule out profitable deviations for the uninformed dealer is

$$b_u + \frac{\delta\gamma}{2}\Delta \geq \Delta \iff b_u \geq \Delta\left(1 - \frac{\delta\gamma}{2}\right).$$

The two conditions combined require that $3\delta\gamma\Delta/4 \geq 1$, a contradiction. Therefore, the uninformed type quotes Δ with probability one in any equilibrium.

Given these quoting strategies, let λ_u and λ_h denote the probability that the customer accepts, respectively, a quote equal to Δ and $1 + \Delta$ in period zero. By looking at the incentive-compatibility constraints of the dealers, we can again derive the relevant bounds on equilibrium values of λ_u and λ_h. Let us start by checking the condition ensuring that the uninformed dealer does not benefit from deviating to $1 + \Delta$. If the uninformed dealer deviates to $1 + \Delta$, her optimal strategy in period one is as follows. If the quote is accepted and the period-one customer is a buyer, then she can trade with probability $1/2$ at a quote of $1 + \Delta$, and this will be her optimal continuation strategy. If the quote is accepted and the period-one customer is a seller, then she cannot trade at a profit (since the dealer's willingness to pay is 0 which is strictly lower than the seller's reservation price equal to $1 - \Delta$), and thus her continuation payoff in this case is zero. If the quote is rejected, our construction of equilibrium in the continuation game following no trade ensures that her continuation payoff is $\frac{\lambda_s^u(nt)+\lambda_b^u(nt)}{4}\Delta$. Hence, her incentive-compatibility constraint can be written as

$$\lambda_u\Delta + \frac{\delta\gamma}{4}(\lambda_u(\lambda_s^u(b_u) + \lambda_b^u(b_u)) + (1-\lambda_u)(\lambda_s^u(nt) + \lambda_b^u(nt)))\Delta \geq \\ \lambda_h(1+\Delta) + \frac{\delta\gamma}{4}(\lambda_h(1+\Delta) + (1-\lambda_h)(\lambda_s^u(nt) + \lambda_b^u(nt))\Delta). \qquad (\mathrm{IC}_{u,\mathrm{post}}(1+\Delta))$$

Next, we check that the uninformed dealer does not benefit from mimicking the low type. If the uninformed dealer deviates to $-1+\Delta$ in period zero, the customer accepts with probability one, and the dealer sells at a loss of $1-\Delta$. If the period-one customer is a buyer, the deviator cannot trade at a profit: Any quote the customer is willing to accept is at most $-1+\Delta$, strictly below the deviator's estimate of the asset value, so her continuation payoff in this case is zero. If the period-one customer is a seller and contacts the deviator (which happens with probability 1/4), the deviator buys the asset at the complete-information price $-1-\Delta$, which the seller accepts with probability one, for an expected gain of $0-(-1-\Delta) = 1+\Delta$. The resulting incentive constraint is

$$\lambda_u\Delta + \frac{\delta\gamma}{4}\Big[\lambda_u\big(\lambda_s^u(b_u) + \lambda_b^u(b_u)\big) + (1-\lambda_u)\big(\lambda_s^u(nt) + \lambda_b^u(nt)\big)\Big]\Delta \;\geq\; \Delta - 1 + \frac{\delta\gamma}{4}(1+\Delta). \qquad (\mathrm{IC}_{u,\mathrm{post}}(-1+\Delta))$$

Observe that the right-hand side of $(\mathrm{IC}_{u,\mathrm{post}}(-1+\Delta))$ is strictly positive if and only if $\delta\gamma > 4(1-\Delta)/(1+\Delta)$, which is compatible with Assumption (2.1); in that region, the constraint requires $\lambda_u \geq \big(\Delta - 1 + \frac{\delta\gamma}{4}(1+\Delta)\big)/\Delta \equiv \lambda_u^* > 0$ in every equilibrium, so the uninformed dealer must trade with strictly positive probability.

Next, we check the incentive constraint of the low-type dealer. If she deviates to Δ, her continuation payoff is $\delta\gamma(1+\lambda_s^l(b_u))\Delta/4$ if the quote is accepted, and $\delta\gamma(1+\lambda_s^l(nt))\Delta/4$ if the quote is rejected. The resulting incentive constraint is

$$\Delta + \frac{\delta\gamma}{2}\Delta \geq \lambda_u(1+\Delta) + \frac{\delta\gamma}{4}(\lambda_u(1+\lambda_s^l(b_u)) + (1-\lambda_u)(1+\lambda_s^l(nt)))\Delta. \qquad (\mathrm{IC}_{l,\mathrm{post}}(\Delta))$$

Finally, we check that the high-type dealer does not benefit from mimicking the uninformed type. In equilibrium, the high type quotes $1+\Delta$, which is accepted with probability λ_h; conditional on trading, the recorded transaction reveals that $v=1$ and the continuation game has complete information, yielding her an expected period-one profit of Δ whenever she is contacted; conditional on not trading, she reaches the no-trade node, where she sells at $1+\Delta$ with probability $\lambda_b^h(nt)$ if the period-one customer is a buyer, and buys at $1-\Delta$ with probability one if the period-one customer is a seller. If she instead deviates to Δ, the customer accepts with probability λ_u, and trade occurs at a loss of $1-\Delta$, and her continuation payoff coincides with her on-path continuation payoff in the event that an uninformed dealer gets contacted in period zero. The resulting incentive constraint is

$$\lambda_h\Delta+\frac{\delta\gamma}{4}\Big[2\lambda_h\Delta+(1-\lambda_h)\big(1+\lambda_b^h(nt)\big)\Delta\Big] \;\geq\; \lambda_u(\Delta-1)+\frac{\delta\gamma}{4}\Big[\lambda_u\big(1+\lambda_b^h(b_u)\big)\Delta+(1-\lambda_u)\big(1+\lambda_b^h(nt)\big)\Delta\Big]. \tag{IC$_{h,\text{post}}(\Delta)$}$$

The above conditions are also sufficient for an equilibrium if we set customers' off-path beliefs equal to the most pessimistic opaque-market beliefs as defined in the proof of Proposition 1, with beliefs following an off-path recorded transaction price specified analogously. Moreover, the conditions are jointly satisfied if we set $\lambda_j^\sigma(nt)=\lambda_j^\sigma(b_u)=0$ for all $\sigma\in\{l,u,h\}$ and $j\in\{b,s\}$, $\lambda_h=0$, and

$$\lambda_u=\max\Big\{0,\ \tfrac{1}{\Delta}\big(\Delta-1+\tfrac{\delta\gamma}{4}(1+\Delta)\big)\Big\}\equiv\lambda_u^*. \tag{A.2}$$

At these values, ($\text{IC}_{u,\text{post}}(-1+\Delta)$) holds with equality whenever its right-hand side is positive, ($\text{IC}_{h,\text{post}}(\Delta)$) and ($\text{IC}_{u,\text{post}}(1+\Delta)$) hold trivially, and ($\text{IC}_{l,\text{post}}(\Delta)$) holds because $\lambda_u\leq\Delta(1+\delta\gamma/4)/(1+\Delta)$, which follows from $\delta\gamma\leq 1<4/(1+2\Delta)$. This establishes equilibrium existence (see Lemma 2). □

A.3.3 Completion of the proof of Theorem 1

Now, we apply the results from the preceding two subsections to establish the welfare ranking in Theorem 1. In the least efficient equilibrium of the opaque market, there is a transaction if and only if the low type is contacted. In the least efficient equilibrium of the post-trade transparent market, λ_h and all continuation acceptance probabilities equal zero, while $\lambda_u=\lambda_u^*$ by ($\text{IC}_{u,\text{post}}(-1+\Delta)$). If $\delta\gamma\leq 4(1-\Delta)/(1+\Delta)$, only the low type trades in period zero and

$$\underline{\Lambda}_{\text{post}}=\alpha\frac{q}{2}+(1-\alpha)\Big(1+\frac{1}{2}(1-q)+\frac{1+q}{4}\Big)\frac{q}{2}>\frac{q}{2}=\underline{\Lambda}_{\text{opaq}}.$$

If $\delta\gamma>4(1-\Delta)/(1+\Delta)$, the uninformed dealer trades with probability λ_u^* in period zero, so $\underline{\Lambda}_{\text{post}}$ increases by $\alpha(1-q)\lambda_u^*$ and hence $\underline{\Lambda}_{\text{post}}>\underline{\Lambda}_{\text{opaq}}$.

To find the most efficient equilibrium in the opaque market, we can apply the welfare bounds derived in Proposition 1 to write:

$$\bar{\Lambda}_{\text{opaq}}=\frac{q}{2}+(1-q)\frac{\Delta}{1+\Delta}+\frac{q}{2}\Big(\frac{\Delta}{1+\Delta}\Big)^2.$$

Next, we consider the upper bound on welfare in the post-trade transparent regime. Denoting $\lambda=(\lambda_u,\lambda_h,\lambda_b^u(b_u),\lambda_s^u(b_u),\lambda_b^h(b_u),\lambda_s^l(b_u),\lambda_b^u(nt),\lambda_s^u(nt),\lambda_b^h(nt),\lambda_s^l(nt))$, the upper bound solves (rescaling the objective by $(1+\delta\gamma)$ to simplify the expression):

$$\begin{aligned}
\max_{\lambda\in[0,1]^{10}} \quad & \left(\frac{q}{2}(1+\lambda_h)+(1-q)\lambda_u\right)+\delta\gamma\Bigg\{\frac{q}{2}(1+\lambda_h) \\
&+\frac{q}{2}(1-\lambda_h)\frac{(1+q)(\lambda_b^h(nt)+1)+(1-q)(\lambda_b^u(nt)+\lambda_s^u(nt))}{4} \\
&+\frac{1-q}{4}\lambda_u\left(\frac{q}{2}(1+\lambda_b^h(b_u))+(2-q)\lambda_b^u(b_u)+\frac{q}{2}(1+\lambda_s^l(b_u))+(2-q)\lambda_s^u(b_u)\right) \\
&+\frac{1-q}{4}(1-\lambda_u)\left(\frac{q}{2}(1+\lambda_b^h(nt))+(2-q)\lambda_b^u(nt)+\frac{q}{2}(1+\lambda_s^l(nt))+(2-q)\lambda_s^u(nt)\right)\Bigg\}
\end{aligned} \tag{Opt$_{\text{post}}$}$$

$$\begin{aligned}
\text{subject to} \quad & (\text{IC}_{l,\text{post}}(\Delta)), (\text{IC}_{u,\text{post}}(1+\Delta)), (\text{IC}_{u,\text{post}}(-1+\Delta)), (\text{IC}_{h,\text{post}}(\Delta)) \\
& \lambda_i^u(j)\le\frac{\Delta}{1+\Delta}, \quad j\in\{b_u,nt\},\ i\in\{b,s\}, \qquad \text{(ICu1)} \\
& \lambda_i^\sigma(j)\le\frac{\Delta\lambda_i^u(j)}{1+\Delta}, \quad j\in\{b_u,nt\},\ (\sigma,i)\in\{(h,b),(l,s)\}. \qquad \text{(IC}\sigma\text{1)}
\end{aligned}$$

Without deriving an explicit solution to the above program, we can obtain a lower bound on $\bar\Lambda_{\text{post}}$ which is strictly higher than $\bar\Lambda_{\text{opaq}}$. To do so, let $\tilde\Lambda_{\text{post}}$ be the value of the objective when we set the vector of control variables λ to be such that $\lambda_u=\Delta/(1+\Delta)$, and ($\text{IC}_{u,\text{post}}(1+\Delta)$), (ICu1), and (IC$\sigma$1) are set to hold with equality. Since this is a feasible value of λ, it follows that $\tilde\Lambda_{\text{post}}\le\bar\Lambda_{\text{post}}$. Moreover, taking the difference with respect to the opaque market, one can verify that the assumption $q\le 2\Delta^2/(1+\Delta^2)$ implies that $\tilde\Lambda_{\text{post}}-\bar\Lambda_{\text{opaq}}>0$, and therefore $\bar\Lambda_{\text{post}}-\bar\Lambda_{\text{opaq}}>0$.

Finally, we prove that the best equilibrium in the opaque market has strictly higher welfare than the worst equilibrium in the post-trade transparent market. Let $x\equiv\Delta/(1+\Delta)$ and $r\equiv\delta\gamma$. In the worst equilibrium under post-trade transparency,

$$\underline{\Lambda}_{\text{post}}=\alpha\frac{q}{2}+\alpha(1-q)\lambda_u^*+(1-\alpha)\frac{7q-q^2}{8},$$

and hence

$$\overline{\Lambda}_{\text{opaq}}-\underline{\Lambda}_{\text{post}}=(1-q)x+\frac{q}{2}x^2-\alpha(1-q)\lambda_u^*-(1-\alpha)\frac{q(3-q)}{8}.$$

Moreover, whenever $\lambda_u^* > 0$,

$$x - \lambda_u^* = \frac{4 - r(1+\Delta)^2}{4\Delta(1+\Delta)} > 0,$$

where the inequality follows from $r \leq 1$ and $\Delta < 1$; since also $x > 0$, we conclude that $\lambda_u^* < x$. Multiplying the welfare difference by $(1+r) > 0$ and using this bound,

$$(1+r)\big(\overline{\Lambda}_{\text{opaq}} - \underline{\Lambda}_{\text{post}}\big) > r\,\Phi + \frac{q}{2}x^2, \qquad \Phi \equiv (1-q)x + \frac{q}{2}x^2 - \frac{q(3-q)}{8}.$$

If $\Phi \geq 0$, the right-hand side is strictly positive. If $\Phi < 0$, then, since $r \leq 1$,

$$r\,\Phi + \frac{q}{2}x^2 \;\geq\; \Phi + \frac{q}{2}x^2 = (1-q)x + qx^2 - \frac{q(3-q)}{8} \;\equiv\; H(x,q).$$

H is strictly increasing in x. Moreover, the assumption $q < 2\Delta^2/(1+\Delta^2)$ is equivalent to $\Delta > c \equiv \sqrt{q/(2-q)}$, and hence to $x > c/(1+c)$. Substituting $q = 2c^2/(1+c^2)$ and $x = c/(1+c)$ into H and clearing the (positive) denominators,

$$8\,H\Big(\tfrac{c}{1+c},\, \tfrac{2c^2}{1+c^2}\Big) \cdot \frac{(1+c^2)^2(1+c)^2}{2c} = (1-c)\Big[4 + 5c - c^2 + 3c^3(1-c)\Big] > 0$$

for all $c \in (0,1)$, since each factor is strictly positive. Hence $H(x,q) > H\big(\tfrac{c}{1+c}, q\big) > 0$, and therefore $\underline{\Lambda}_{\text{post}} < \overline{\Lambda}_{\text{opaq}}$.

A.4 Proof of Theorem 2

Throughout, we rely on results proven in Appendix A.1. We begin by characterizing the structure of equilibria in the pre-trade transparent market.

A.4.1 The pre-trade transparent market

Proposition 3 (Structure of equilibria in the pre-trade transparent market)**.** *In all equilibria, when $Z_t = 0$, the high type quotes 1, the uninformed type quotes a deterministic price $b_u \leq \Delta$, and the low type quotes a random price distributed on $[\underline{b}, -1+\Delta]$, with $\underline{b} > -1$. Given a pair of on-path quotes (q_A, q_B), the customer accepts the lower quote with probability one; if $q_A = q_B$, the customer breaks the tie between dealers uniformly at random but refuses to trade with probability $1 - \lambda_u$ when $q_A = q_B = b_u$.*

When $Z_t = 1$, trade in period t is fully efficient.

There exists a fully efficient equilibrium (with $\lambda_u = 1$) if and only if $\delta\gamma \leq \bar{\delta}_{\text{pre}}$, where

$\bar{\delta}_{\text{pre}} > 0$ *if and only if* $q < 2\Delta - 1$.

Proof. Step 1: Continuation equilibrium in period one. Consider the game that ensues in period one following any pair of on-path quotes in period zero. Because under the pre-trade transparent regime dealers observe each other's quotes at the end of each period, both dealers are symmetrically informed and hold a common belief about the asset's value at the time of quoting. The period-one customer, on the other hand, enters the market with beliefs about v equal to the prior.

Fix an on-path pair of period-zero quotes, and let (σ_i, σ_{-i}) denote the corresponding period-zero types of the two dealers. Since period-zero strategies are separating and dealers observe both quotes, (σ_i, σ_{-i}) is common knowledge among the dealers in period one. Let $v' \in \{-1, 0, 1\}$ denote their common posterior mean of the asset value. Thus, if $v' = -1$, then $\sigma_i, \sigma_{-i} \in \{l, u\}$ and at least one dealer has type l; if $v' = 0$, then $\sigma_i = \sigma_{-i} = u$; and if $v' = 1$, then $\sigma_i, \sigma_{-i} \in \{u, h\}$ and at least one dealer has type h. For each initial type σ that can arise together with common belief v', let $\bar{b}(\sigma, v')$ denote the supremum of the support of that dealer's period-one quoting strategy.

We first show that no dealer can quote above $v' + \Delta$. Suppose by contradiction, that $\bar{b}(\sigma_i, -1) > -1+\Delta$ for some $\sigma_i \in \{l, u\}$. Since such a quote is accepted with zero probability, dealer i's payoff is zero. This implies that $\bar{b}(\sigma_{-i}, -1) = -1$; for if not dealer i could attain a strictly positive payoff by deviating to $b \in (-1, \bar{b}(\sigma_{-i}, -1))$. Hence, dealer $-i$'s payoff is zero. But then dealer $-i$ can profitably deviate to $b \in (-1, -1 + \Delta)$ which leads to trade at a profit with strictly positive probability. This is a contradiction.

The same argument implies that $\bar{b}(u, 0) \leq \Delta$. Indeed, suppose instead that $\bar{b}(u, 0) > \Delta$, in which case dealers' payoff in this part of the game is zero. Consider the uninformed dealer's deviation to a quote $\tilde{b} \in (0, \Delta)$, sufficiently close to zero. Before observing period-one quotes, the customer's belief about the dealers' posterior means is supported on $\{(-1, -1), (0, 0), (1, 1)\}$. Hence, after observing a pair $(b_u, \tilde{b})$, where b_u is a quote above Δ in the support of the uninformed dealer's strategy when $v' = 0$, the never-dissuaded-once-convinced and unprejudiced-belief refinements imply that the customer must continue to assign posterior mean 0 to the asset. She is therefore willing to accept $\tilde{b}$ whenever it is the lowest quote. This gives the deviating dealer strictly higher payoff, a contradiction. An analogous argument, using the fact that quotes associated with common beliefs below 1 are bounded away from $1 + \Delta$, shows that $\bar{b}(\sigma, 1) \leq 1 + \Delta$ for every $\sigma \in \{u, h\}$ that can arise when $v' = 1$.

We now show that in fact dealers never quote a price higher than v'. Fix $v' \in \{-1, 0, 1\}$ and define $\bar{b}(v') \equiv \max\{\bar{b}(\sigma_i, v'), \bar{b}(\sigma_{-i}, v')\}$. Suppose, toward a contradiction, that $\bar{b}(v') > v'$. Without loss of generality, suppose that $\bar{b}(v') = \bar{b}(\sigma_i, v')$. Then, dealer i's payoff in this part of the game is zero. This implies that $\bar{b}(\sigma_{-i}, v') = v'$ for if not dealer i would strictly benefit from quoting $\tilde{b} \in (v', \bar{b}(\sigma_{-i}, v'))$ which, given that the competitor's quote pins down the customer's belief, leads to trade with positive probability at a strictly profitable price. This in turn implies that dealer j's payoff must be zero as well. But then, dealer j can profitably deviate to $\tilde{b}' \in (v', \bar{b}(v'))$: Since $\bar{b}(v') \leq v' + \Delta$ and the customer's posterior mean after observing the competitor's quote is v', this quote is accepted with positive probability. This is a contradiction. Thus, $\bar{b}(v') \leq v'$.

Finally, any quote below v' yields a negative payoff. Hence, both dealers quote v' with probability one. Given this quoting strategy, the customer strictly prefers to trade after observing (v', v'), and hence trade occurs with probability one. Any off-path beliefs for the period-one customer can sustain this continuation equilibrium, since dealers can only weakly lose from deviating to any other quote.

<u>Step 2: Equilibrium strategies in period zero.</u> First, we derive necessary conditions on dealers' period-zero quoting strategies. We then show that these conditions imply that there are two kinds of equilibria that may arise in period zero: a fully efficient one, and one where the uninformed type trades with interior probability. The latter equilibrium always exists under Assumption (2.1), whereas the former one exists if and only if $\delta\gamma \leq \bar{\delta}_{\text{pre}}$ for some $\bar{\delta}_{\text{pre}} < 1$ that we derive explicitly in the proof below.

2.a) Necessary conditions for equilibrium strategies. Let $\underline{b}_\sigma$ and $\bar{b}_\sigma$ be the infimum and supremum of the quoting strategy of a dealer with type $\sigma \in \{l, u, h\}$ in period zero. First, we show that $\underline{b}_h = \bar{b}_h = 1$. Suppose that $\bar{b}_h > 1 + \Delta$, in which case the equilibrium payoff of the high type equals zero. Then, under the unprejudiced-belief refinement, the high type can do strictly better by quoting $b \in (1, 1 + \Delta)$, which leads to profitable trade in the event that the competitor has a high type. Thus, $\bar{b}_h \leq 1 + \Delta$. Next, suppose that $\bar{b}_h \in (1, 1 + \Delta)$. Then, an analogous argument implies that the high type would strictly prefer to quote $\bar{b}_h - \epsilon$, with $\epsilon > 0$ sufficiently small, a contradiction. Hence, $\bar{b}_h \leq 1$. Since $\underline{b}_h \geq 1$, it follows that $\underline{b}_h = \bar{b}_h = 1$.

For the low and the uninformed types, the continuation payoff if they play their equilibrium strategy in period zero is also 0. Since this is the lowest payoff they could get, it follows that any deviation would lead to a weak improvement in terms of continuation payoffs. Hence, we use the fact that a necessary condition for these two types to not deviate is that flow payoffs in period zero are maximized by any quote in the support of their equilibrium strategy.

Let us now derive the quoting strategy of the low type. We start by noting that $\underline{b}_l > -1$. Otherwise, the low type makes at most a zero payoff, which is a contradiction since she can do strictly better by quoting a price slightly below $-1 + \Delta$ and trading with probability at least $1-q$ (since $\underline{b}_u \geq 0 > -1+\Delta$ and the customer accepts the lowest quote with probability one whenever that quote is below $-1+\Delta$). This implies that the low type's equilibrium flow payoff in period zero is strictly positive, and that her equilibrium strategy must be atomless (otherwise, undercutting the price that occurs with strictly positive probability would be a profitable deviation). Also, $\bar{b}_l \leq -1+\Delta$; otherwise, the low type would get a zero payoff by quoting $\bar{b}_l$ (which is rejected by the customer for sure). Moreover, it must be that $\bar{b}_l \geq -1+\Delta$. This is because, if $\bar{b}_l < -1+\Delta$, the low type's payoff from quoting a price close to $\bar{b}_l$ is approximately $(1-q)(\bar{b}_l+1)$. This can be improved upon by quoting a price close to $-1+\Delta < \underline{b}_u$, which yields a flow payoff of $(1-q)\Delta$. Hence, $\bar{b}_l = -1+\Delta$ and the low type's period-zero payoff is $(1-q)\Delta$. In order to make her indifferent across all quotes in $[\underline{b}_l, -1+\Delta]$, we require that

$$(q(1-F_l(b))+1-q)(b+1) = (1-q)\Delta, \quad \forall b \in [\underline{b}_l, -1+\Delta],$$

where $F_l(b)$ is the cumulative distribution function of the low type's quoting strategy. This yields $F_l(b) = 1 - \frac{(1-q)(-1+\Delta-b)}{q(1+b)}$ and $\underline{b}_l = (1-q)(-1+\Delta) - q$.

Let us now consider the uninformed dealer. First, suppose that $\bar{b}_u > 1$. Such a quote must be rejected by the customer with probability one; otherwise, the high-type dealer, whose equilibrium payoff in period zero is 0, would benefit from deviating to $\bar{b}_u$. But then the uninformed dealer can do strictly better by quoting 1 which leads to trade at a profit with strictly positive probability (when the competitor is uninformed and quotes a price close to $\bar{b}_u$). Next, let us show that $\bar{b}_u \leq \Delta$. If $\Delta < \bar{b}_u < 1$, then the uninformed dealer who quotes $\bar{b}_u$ only gets to trade if the competitor's type is high. This is because either the uninformed competitor quotes a price strictly lower than $\bar{b}_u$ (which is preferred by the customer), or she quotes exactly $\bar{b}_u$ (in which case the customer will refuse to trade). But then, the expected payoff from quoting $\bar{b}_u$ is $\frac{q}{2}(\bar{b}_u - 1) < 0$, a contradiction.

In the same way, in order to rule out the uninformed type making a strictly negative payoff, we need that the uninformed type's strategy has an atom of size $\beta > 0$ at $\bar{b}_u$ and that $\bar{b}_u > 0$ (otherwise, the uninformed type makes losses with this quote due to adverse selection). We now argue that this implies that $\underline{b}_u = \bar{b}_u$. The uninformed dealer's payoff from quoting $\bar{b}_u$ is $(1-q)\frac{\beta}{2}\bar{b}_u + \frac{q}{2}(\bar{b}_u - 1)$. Suppose that $\underline{b}_u < \bar{b}_u$ and consider the uninformed dealer's payoff from quoting $b \in (\underline{b}_u, \bar{b}_u)$. Since beliefs are monotone, the customer's posterior mean about v after observing this quote is at least 0, and since $b < \Delta$ the customer will be willing to accept with probability one whenever b is the lowest quote she observes. Hence, the payoff from $b \approx \bar{b}_u$ is approximately $(1-q)\beta\bar{b}_u + \frac{q}{2}(\bar{b}_u - 1)$ which is strictly higher than her equilibrium payoff. Given this, let $b_u = \underline{b}_u = \bar{b}_u \leq \Delta$.

2.b) Construction of quoting strategies in a fully efficient equilibrium. Suppose first that $b_u < \Delta$. Observe that, in equilibrium, with probability one, the customer strictly benefits from accepting the lowest quote (since the quote pairs in which the customer is indifferent arise with probability zero). Hence, the outcome in such an equilibrium will be fully efficient, and this will be the most efficient equilibrium, provided that it exists.

Let us check dealers' incentive constraints in period zero. First, consider the uninformed dealer's incentive to deviate to a quote of 1.[41] After doing so, her continuation payoff if the competitor is informed is zero. If the competitor is uninformed and the future customer is a buyer, she can quote 1 and trade with probability 1/2, or quote Δ and trade with probability one (either one of these strategies can be optimal depending on parameter values, and one of them yields her optimal continuation payoff under the most punitive beliefs by the future customer). Otherwise, if the future customer is a seller, she cannot trade at a profit and thus her continuation payoff is zero. This yields the constraint

$$\frac{1-q}{2}b_u + \frac{q}{2}(b_u - 1) \geq \delta\gamma\frac{1-q}{4}\kappa \quad \iff \; b_u \geq q + \delta\gamma(1-q)\kappa/2 \equiv \underline{b}_u(\text{pre}),$$

where $\kappa = \max\{1, 2\Delta\}$ captures the two possibilities for continuation play in period one. Note that this condition implies that the uninformed dealers makes non-negative profits in equilibrium.

The low type's equilibrium payoff in period zero is $(1-q)\Delta$. If she deviates to b_u, she gets $\frac{1-q}{2}(b_u + 1)$ in period zero. If the competitor is uninformed and the future customer is a buyer, then in period one, the low type may either quote 0 and trade with probability 1/2 or quote $-1 + \Delta$ and trade with probability one. Continuation payoffs conditional on the future customer being a seller and the competitor being uninformed are zero. As a result, the constraint preventing this deviation is

$$(1-q)\Delta \geq (1-q)\frac{b_u + 1}{2} + \delta\gamma\frac{1-q}{4}\kappa \iff b_u \leq 2\Delta - 1 - \delta\gamma\kappa/2 \equiv \bar{b}_u(\text{pre}) < \Delta.$$

Hence, for a fully efficient equilibrium to exist, we need:

$$\underline{b}_u(\text{pre}) \leq \bar{b}_u(\text{pre}) \iff \delta\gamma \leq \frac{2(2\Delta - 1 - q)}{(2-q)\kappa} \equiv \bar{\delta}_{\text{pre}}. \tag{A.3}$$

As we argue below, this condition is also sufficient.

2.c) Construction of quoting strategies in an inefficient equilibrium. Second, let us consider

[41] By symmetry, the uninformed type's continuation payoff must be the same whether she mimics the low or the high type; however, mimicking the high type is less costly in period zero, so ruling out this deviation also rules out the deviation to a quote of the low type.

the case with $b_u = \Delta$. After observing a pair of quotes equal to (Δ, Δ), the customer is indifferent between accepting either quote (with the tie broken uniformly at random by symmetry of strategies) or rejecting both. Let λ_u denote the probability that she accepts one of the quotes. As before, we have to check that the uninformed and low types do not benefit from deviating to, respectively, 1 and b_u (other deviations can only be less profitable). For the uninformed type, we then obtain the following incentive constraint:

$$\frac{(1-q)\lambda_u}{2}\Delta + \frac{q}{2}(\Delta - 1) \geq \delta\gamma\frac{1-q}{4}\kappa \iff \lambda_u \geq \frac{\delta\gamma(1-q)\kappa/2 + q(1-\Delta)}{(1-q)\Delta} \equiv \underline{\lambda}_u^1(\text{pre}).$$

The incentive constraint for the low type is

$$(1-q)\Delta \geq \frac{(1-q)\lambda_u}{2}(\Delta+1) + \delta\gamma\frac{1-q}{4}\kappa \iff \lambda_u \leq \frac{2\Delta - \delta\gamma\kappa/2}{1+\Delta} \equiv \bar{\lambda}_u(\text{pre}).$$

Additionally, to rule out the high type deviating to Δ, it must hold that

$$0 \geq \left[\frac{(1-q)\lambda_u}{2} + q\right](\Delta - 1) + \delta\gamma\frac{1-q}{4}\kappa \iff \lambda_u \geq \frac{2q(\Delta-1) + \delta\gamma(1-q)\kappa/2}{(1-q)(1-\Delta)} \equiv \underline{\lambda}_u^2(\text{pre}).$$

Let $\underline{\lambda}_u(\text{pre}) \equiv \max\{\underline{\lambda}_u^1(\text{pre}), \underline{\lambda}_u^2(\text{pre})\}$. For this equilibrium to exist, we therefore need

$$\underline{\lambda}_u(\text{pre}) \leq \bar{\lambda}_u(\text{pre}) \iff \delta\gamma \leq \min\left\{\frac{2[(2-q)\Delta^2 - q]}{(1-q)(1+2\Delta)\kappa}, \frac{2(1-\Delta)(\Delta+q)}{\kappa(1-q)}\right\}. \tag{A.4}$$

Assumption (2.1) implies that the necessary condition for existence of an inefficient equilibrium (A.4) is satisfied.

2.d) Construction of off-path beliefs. For both kinds of equilibria, the above conditions are also sufficient if we specify the customer's beliefs as follows: After any off-path pair of quotes, she assigns probability one to the profile of dealers' estimates minimizing the implied posterior mean of the asset value within the set permitted by the applicable refinements, and accepts a quote if and only if trading at that quote is weakly profitable given her posterior. In period zero, this rule gives: After (b_u, b) with $b < b_u$, the unprejudiced beliefs refinement attributes the pair to an uninformed dealer quoting b_u and a low-type quoting b, so the customer rejects; after $(1, b)$ with $b \in (b_u, 1)$, it attributes the pair to a high type and an uninformed deviator, so the customer accepts. In period one, where the customer's prior is supported on profiles with a common estimate: After $(0, b)$ with $b < 0$, both estimates equal -1 (the interpretation that both quotes reflect a low asset value), so the customer accepts only if $b \leq -1 + \Delta$; after $(0, b)$ with $b > 0$, she believes that both dealers are uninformed and accepts the quote of 0. It is straightforward to verify that no dealer benefits from deviating to an off-path quote under these beliefs. Thus, an equilibrium exists in this regime (see Lemma 2). □

A.4.2 Completion of the proof of Theorem 2

Let us start by showing the first part of the theorem involving weak welfare dominance. Among all equilibria in the pre-trade transparent regime that were identified in Proposition 3, the least efficient one is the one where, in period zero, the uninformed type quotes Δ and the customer agrees to trade after observing a pair of quotes (Δ, Δ) with probability $\underline{\lambda}_u(\text{pre})$. The outcome is fully efficient in all other contingencies. This inefficient equilibrium always exists under Assumption (2.1). As a result, the worst-case welfare under this protocol is

$$\underline{\Lambda}_{\text{pre}} = 1 - \alpha(1-q)^2(1 - \underline{\lambda}_u(\text{pre})).$$

To compare it to the welfare in the least efficient equilibrium in the post-trade transparent regime, note that when $Z = 1$, we have $\underline{\Lambda}_{\text{pre}}(1) = 1 > \underline{\Lambda}_{\text{post}}(1)$. When $Z = 0$,

$$\underline{\Lambda}_{\text{pre}}(0) = 1 - (1-q)^2\big(1 - \underline{\lambda}_u(\text{pre})\big), \qquad \underline{\Lambda}_{\text{post}}(0) = \frac{q}{2} + (1-q)\lambda_u^*,$$

where λ_u^* is defined in equation (A.2)

We first note that $\underline{\lambda}_u(\text{pre}) > \lambda_u^*$. This is immediate if $\lambda_u^* = 0$. If $\lambda_u^* > 0$, then $\delta\gamma > 4(1-\Delta)/(1+\Delta)$; since $\delta\gamma \leq 1$, this implies $\Delta > 3/5$ and hence κ—which was defined as $\max\{1, 2\Delta\}$—equals 2Δ. Therefore,

$$\underline{\lambda}_u(\text{pre}) \geq \underline{\lambda}_u^1(\text{pre}) = \delta\gamma + \frac{q(1-\Delta)}{(1-q)\Delta},$$

using the definitions from the proof of Proposition 3, and

$$\underline{\lambda}_u^1(\text{pre}) - \lambda_u^* = \frac{\delta\gamma(3\Delta - 1)}{4\Delta} + \frac{1-\Delta}{(1-q)\Delta} > 0.$$

It follows that

$$\begin{aligned}\underline{\Lambda}_{\text{pre}}(0) - \underline{\Lambda}_{\text{post}}(0) &= \frac{3q}{2} - q^2 + (1-q)^2\underline{\lambda}_u(\text{pre}) - (1-q)\lambda_u^* \\ &> \frac{3q}{2} - q^2 - q(1-q)\lambda_u^* \geq \frac{q}{2} > 0,\end{aligned}$$

where the weak inequality uses $\lambda_u^* \leq 1$.

Next, we compare the regimes in terms of their most efficient equilibrium. If condition

(A.3) is satisfied, then the pre-trade transparent market is fully efficient ($\bar{\Lambda}_{\text{pre}} = 1$) and the result follows immediately. Suppose that (A.3) is violated. Then, the most efficient equilibrium of the pre-trade transparent market is the one where the uninformed dealer quotes Δ in period zero and, conditional on a pair of quotes (Δ, Δ), the customer agrees to trade with probability $\bar{\lambda}_u(\text{pre})$.

First, $\bar{\Lambda}_{\text{pre}}(1) = 1 > \bar{\Lambda}_{\text{post}}(1)$. To compare the regimes conditional on $Z = 0$, we write

$$\bar{\Lambda}_{\text{pre}}(0) = 1 - (1-q)^2(1 - \bar{\lambda}_u(\text{pre})).$$

Moreover, letting $\bar{\lambda}_u(\text{post})$ and $\bar{\lambda}_h(\text{post})$ be part of a solution to the program (Opt_{post}), we can write

$$\begin{aligned}\bar{\Lambda}_{\text{post}}(0) = 1 - \frac{q}{2}(1 - \bar{\lambda}_h(\text{post})) - (1-q)(1 - \bar{\lambda}_u(\text{post}))\\ \leq 1 - (1-q)(1 - \bar{\lambda}_u(\text{post})).\end{aligned}$$

Thus, it suffices to show that $(1-q)(1-\bar{\lambda}_u(\text{pre})) \leq 1 - \bar{\lambda}_u(\text{post})$. To do this, write

$$1 - \bar{\lambda}_u(\text{post}) - (1-q)(1 - \bar{\lambda}_u(\text{pre})) \geq \frac{1 - \delta\gamma\Delta/4}{1+\Delta} - \frac{(1-q)(1 - \Delta + \delta\gamma\kappa/2)}{1+\Delta} > 0,$$

where the first inequality follows from the fact that $\bar{\lambda}_u(\text{post}) \leq \Delta(1 + \delta\gamma/4)/(1+\Delta)$ by ($\text{IC}_{l,\text{post}}(\Delta)$) and the final inequality follows from the assumption that $\delta\gamma \leq \bar{\delta}_1$ (see Lemma 2).

Finally, we establish the second part of Theorem 2. Suppose that $q \geq \underline{q} \equiv \frac{4\Delta^2 - \delta\gamma(2\kappa(1+\Delta) - \Delta^2)}{2(2-\delta\gamma\kappa)(1+\Delta)}$. When $Z = 1$, we have $\underline{\Lambda}_{\text{pre}}(1) = 1 > \bar{\Lambda}_{\text{post}}(1)$. When $Z = 0$, we showed above that

$$\bar{\Lambda}_{\text{post}}(0) \leq 1 - (1-q)(1 - \bar{\lambda}_u(\text{post})) \leq 1 - \frac{(1-q)(1 - \delta\gamma\Delta/4)}{1+\Delta} \leq 1 - (1-q)^2(1 - \underline{\lambda}_u(\text{pre})),$$

where the final inequality follows from the assumption that $q \geq \underline{q}$. Hence, $\underline{\Lambda}_{\text{pre}} \geq \bar{\Lambda}_{\text{post}}$ when q is high enough ($\bar{\Lambda}_{\text{pre}} > \underline{\Lambda}_{\text{post}}$ follows from the proof as well).[42]

A.5 Proof of Theorem 3

Throughout, we rely on results proven in Appendix A.1. We begin by characterizing equilibria in the fully transparent regime.

[42] Recall that we assumed that $q < 2\Delta^2/(1+\Delta^2)$ in Section 2.4. One can check that $\underline{q} < 2\Delta^2/(1+\Delta^2)$, and thus the region where $q > \underline{q}$ and $q < 2\Delta^2/(1+\Delta^2)$ is non-empty.

A.5.1 The fully transparent market

Proposition 4 (Structure of equilibria in the fully transparent market)**.** *Equilibria in the fully transparent market have the same structure as in the case of a pre-trade transparent market. However, a fully efficient equilibrium exists if and only if* $\delta\gamma \leq \bar{\delta}_{\text{full}} < \bar{\delta}_{\text{pre}}$*, and, in equilibria in which the uninformed dealers quote* $b_u = \Delta$*, the highest* λ_u *that can be sustained is strictly lower than in the pre-trade transparent market.*

Proof. Step 1: Continuation equilibrium in period one. As in the pre-trade transparent market, it is still the case that dealers learn each other's types by observing the quotes provided in period zero. The only difference, relative to that regime, is that now the period-one customer will enter the market with a belief about v that is refined by the outcome of trade in period zero. Still, the arguments that we gave in the proof of Proposition 3 for showing that the Bertrand outcome is the unique continuation equilibrium candidate in period one apply here as well with the following adjustment: The period-one customer's beliefs μ^c are now derived from the publicly disclosed transaction record via Bayes' rule, and the arguments apply with $\text{supp}(\mu^c)$ in place of the full set of common-estimate profiles.
Step 2: Equilibrium strategies in period zero.
2.a) Necessary conditions for equilibrium strategies. Applying the arguments from Proposition 3, it must be that the structure of dealers' strategies coincides with the one in the pre-trade transparent market. Namely, the low and high-type dealers will follow exactly the same quoting strategy as in the pre-trade transparent regime, and the uninformed dealer will play a pure strategy of quoting $b_u \leq \Delta$. However, as we show next, the introduction of post-trade transparency will affect dealers' continuation payoffs conditional on deviating in period zero, and hence will have an effect on trading outcomes through dealers' incentive-compatibility constraints. More precisely, the equilibrium values of b_u and the probability that the customer agrees to trade when observing a pair of quotes equal to (Δ, Δ) will be different relative to the regime without post-trade transparency.
2.b) Construction of a fully efficient equilibrium. The condition preventing the uninformed type from deviating to 1 is unchanged (recall that future traders only observe the transaction price, not individual quotes), and hence we must have $b_u \geq \underline{b}_u(\text{pre})$. Consider the low-type dealer. Her on-path payoffs are the same as those in the pre-trade transparent market, so we only need to compute her payoff from deviating to b_u in period zero. If the competitor's type is l, the winning quote will reveal that $v = -1$ and the continuation payoff is 0. If the competitor's type is u, there will be trade at a price equal to b_u and hence the period-one customer will enter the market with beliefs that assign probability zero to type l. Because of the never-dissuaded-once-convinced refinement, this in turn implies that, after deviating

in period zero, the low-type dealer can undercut her competitor's quote (which is equal to 0) in the following period if the customer is a buyer without facing belief-based punishment, which gives her a continuation payoff arbitrarily close to 1. If the future customer is a seller, the continuation payoff is 0. As a result, her incentive constraint in period zero becomes

$$(1-q)\Delta \geq \frac{1-q}{2}(b_u+1) + \frac{\delta\gamma}{2}(1-q) \iff b_u \leq 2\Delta - \delta\gamma - 1 \equiv \bar{b}_u(\text{full}).$$

Observe that $\bar{b}_u(\text{full}) < \bar{b}_u(\text{pre})$, which in turn implies that we have a more stringent necessary condition for existence of an efficient equilibrium in the fully transparent market, given by

$$\underline{b}_u(\text{pre}) \leq \bar{b}_u(\text{full}) \iff \delta\gamma \leq \frac{2(2\Delta-1-q)}{(1-q)\kappa+2} \equiv \overline{\delta}_{\text{full}} < \overline{\delta}_{\text{pre}}.$$

2.c) Construction of an inefficient equilibrium. Next, consider equilibria in which, in period zero, the uninformed dealer quotes $b_u = \Delta$ and the customer agrees to trade with probability λ_u after observing a pair of quotes equal to (Δ, Δ). As in the previous case, the incentive-compatibility constraint of the uninformed and the high dealers is the same as in the pre-trade transparent market, which yielded the lower bound $\lambda_u \geq \underline{\lambda}_u(\text{pre})$. For the low-type dealer, the continuation payoff if she deviates to Δ in period zero is again equal to $(1-q)/2$. Thus, we have

$$(1-q)\Delta \geq \frac{(1-q)\lambda_u}{2}(\Delta+1) + \frac{\delta\gamma}{2}(1-q) \iff \lambda_u \leq \frac{2\Delta-\delta\gamma}{1+\Delta} \equiv \bar{\lambda}_u(\text{full}).$$

Observe that $\bar{\lambda}_u(\text{full}) < \bar{\lambda}_u(\text{pre})$ as stated in the proposition.

By the above, an equilibrium in which the uninformed dealer quotes Δ in period zero exists only if

$$\underline{\lambda}_u(\text{pre}) \leq \bar{\lambda}_u(\text{full}) \iff$$
$$\delta\gamma \leq \min\left\{\frac{2[(2-q)\Delta^2-q]}{(1-q)[(1+\Delta)\kappa+2\Delta]}, \frac{4(1-\Delta^2)(\Delta+q)}{(1+\Delta)(1-q)[\kappa(1+\Delta)+2(1-\Delta)]}\right\},$$

where the expressions in the curly bracket correspond to $\overline{\delta}_1$ and $\overline{\delta}_2$, respectively, as defined in Lemma 2.

2.d) Equilibrium existence. The above conditions are also sufficient for an equilibrium if we set off-path beliefs in period zero in the same way as we did in the proof of Proposition 3. Thus, an equilibrium exists if $q < 2\Delta^2/(1+\Delta^2)$ and $\delta\gamma \leq \overline{\delta}_{\text{exist}}$. This final step also concludes the proof of equilibrium existence in all the regimes (Lemma 2). $\square$

A.5.2 Completion of the proof of Theorem 3

First, both regimes share the same least efficient equilibrium, with welfare

$$\underline{\Lambda}_{\text{pre}} = \underline{\Lambda}_{\text{full}} = \alpha(1-(1-q)^2(1-\underline{\lambda}_u(\text{pre}))) + 1 - \alpha,$$

where α is defined in (A.1). Therefore, it suffices to compare the two regimes in terms of their most efficient equilibrium.

If $\delta\gamma \leq \overline{\delta}_{\text{full}} < \overline{\delta}_{\text{pre}}$, then both regimes have a fully efficient equilibrium. If $\delta\gamma \in (\overline{\delta}_{\text{full}}, \overline{\delta}_{\text{pre}}]$, then only the pre-trade transparent market has a fully efficient equilibrium, and hence $\bar{\Lambda}_{\text{pre}} = 1 > \bar{\Lambda}_{\text{full}}$. Otherwise, if $\delta\gamma > \overline{\delta}_{\text{pre}}$, then neither regime has a fully efficient equilibrium. The respective upper bounds on equilibrium welfare are given by

$$\bar{\Lambda}_{\text{pre}} = \alpha(1-(1-q)^2(1-\bar{\lambda}_u(\text{pre}))) + 1 - \alpha, \quad \bar{\Lambda}_{\text{full}} = \alpha(1-(1-q)^2(1-\bar{\lambda}_u(\text{full}))) + 1 - \alpha.$$

Since $\bar{\lambda}_u(\text{pre}) > \bar{\lambda}_u(\text{full})$, it follows that $\bar{\Lambda}_{\text{pre}} > \bar{\Lambda}_{\text{full}}$. Hence, the pre-trade transparent regime weakly welfare-dominates the fully transparent one if and only if $\delta\gamma > \overline{\delta}_{\text{full}}$.

A.6 Proof of Theorem 4

For $j \in \{b, s\}$, let $(x_Z^j(\theta_i, \theta_{-i}), p_Z^j(\theta_i, \theta_{-i}), t_Z^j(\theta_i, \theta_{-i}))$ denote the objects defined in Section 4 conditional on the customer being, respectively, a buyer or a seller. Using the discounted probabilities derived in Lemma 1, the optimal mechanism maximizes

$$\mathbb{E}\left[\sum_{i\in\{A,B\},\, j\in\{b,s\}} (x_0^j(\theta_i, \theta_{-i}) + \delta\gamma x_1^j(\theta_i, \theta_{-i}))\right] \tag{Opt_{mech}}$$

over all feasible mechanisms as defined in Section 4, with only one modification: we require incentive compatibility to hold for every configuration of customer profiles $(j, h) \in \{b, s\}^2$ that may arise in periods zero and one. We show below that this stronger notion of feasibility turns out to be irrelevant and thus it is without loss of generality to define feasibility as in Section 4.

<u>Step 1: Solving for the optimal mechanism.</u> Suppose first that $\Delta \geq 1/2$. In that case, the mechanism proposed in the proposition is fully efficient and hence cannot be improved upon. One can also readily check that it is feasible.

Next, suppose that $\Delta < 1/2$. We can then show that any implementable mechanism satisfies $x_0^j(0,0) + \delta\gamma x_1^j(0,0) \leq (1+\delta\gamma)\Delta$ for all $j \in \{b, s\}$. To do so, consider the case with $j = b$, and combine the ex-post IC constraint of the low type when her competitor is uninformed with the ex-post IR constraint of the uninformed type when her competitor is uninformed to obtain that

$$x_0^b(0,0) + \delta\gamma x_1^b(0,0) \leq x_0^b(-1,0) + p_0^b(-1,0) + \delta\gamma(x_1^b(-1,0) + p_1^b(-1,0)).$$

On the other hand, for all $Z \in \{0,1\}$,

$$\begin{aligned}\Delta &\geq x_Z^b(-1,0) + x_Z^b(0,-1) + t_Z^b(0,-1) \geq x_Z^b(-1,0) + x_Z^b(0,-1) + p_Z^b(-1,0) + p_Z^b(0,-1) \\ &\geq x_Z^b(-1,0) + p_Z^b(-1,0),\end{aligned}$$

where the first inequality comes from buyer's ex-post IR constraint, the second one from the designer's ex-post budget-balance constraint, and the third one from the dealer's ex-post IR constraint, all evaluated at the type profile $(0, -1)$. These two conditions combined imply $x_0^b(0,0) + \delta\gamma x_1^b(0,0) \leq (1+\delta\gamma)\Delta$. The case with $j = s$ can be obtained symmetrically by considering the IC constraint of type $\theta = 1$.

The above implies that welfare under any implementable mechanism is bounded above by $\Lambda^* \equiv 1 - (1-q)^2(1-2\Delta)$, i.e., welfare attained when trade happens with probability one under all contingencies, except when both dealers are uninformed and each trades with probability Δ. Our proposed mechanism is feasible and attains precisely that bound.

Step 2: Comparison with the pre-trade transparent market. The comparison with respect to the most efficient equilibrium in the pre-trade transparent market is trivial when $\Delta \geq 1/2$. If $\Delta < 1/2$, note that $\overline{\delta}_{\text{pre}} < 0$, and thus the pre-trade transparent regime fails to have a fully efficient equilibrium. Thus, welfare in the most efficient equilibrium is given by $\overline{\Lambda}_{\text{pre}} = 1 - \alpha(1-q)^2(1 - \overline{\lambda}_u(\text{pre}))$. The welfare difference is

$$\Lambda^* - \overline{\Lambda}_{\text{pre}} = \frac{(1-q)^2[2\Delta^2 - \delta\gamma(1/2 - \Delta - 2\Delta^2)]}{(1+\Delta)(1+\delta\gamma)} > 0,$$

where the inequality follows from the assumption on $\delta\gamma$ in Lemma 2 (i.e., $\delta\gamma \leq \bar{\delta}_1$).

A.7 Proof of Theorem 5

Step 1: Specifying equilibrium strategies. We construct an ex-post equilibrium of the game where, in period zero, dealer i with type θ quotes θ. In period one, if the customer is a buyer, dealer i with type 0 quotes $b_u^1 \in (0,1)$ if she learned that her competitor has a high type, and quotes $-1+\Delta$ if she learned that her competitor has a low type in the previous period. Strategies are set symmetrically if the period-one customer is a seller. In all other cases, dealer i with type θ quotes θ in period one. The customer agrees to trade with probability one at the end of the auction when she is indifferent in all cases except after she observes that both dealers quoted 0, in which case she accepts with probability $2\min\{\Delta, 1/2\}$ (with the tie broken uniformly across the two dealers) and exits otherwise. The only on-path case in which she strictly prefers to exit or request another quote over accepting is when only one dealer quotes 0, or when only one dealer quotes $-1+\Delta$, which only happens when one dealer is uninformed and the other one has a low type. In such a case, we set her strategy to be to ask for another quote with probability one. Conditional on the customer having requested another quote, the winning dealer's strategy is to counteroffer $\hat{v}+\Delta$, where $\hat{v}$ is the customer's estimate of the value at this stage, and the customer accepts with probability one. These strategies implement the allocation from Theorem 4.

Step 2: Checking players' incentive constraints. At any off-path pair of quotes, except where noted explicitly below, the customer's beliefs assign probability one to the profile of dealers' estimates minimizing the implied posterior mean of the asset value across all profiles; she accepts if and only if trading at the prescribed price is weakly profitable given her posterior.

Let us now check that this is an equilibrium by solving the game backwards. The strategies of the customer and the winning dealer are clearly optimal in the node of the game in which the customer requests another quote. Let us check dealers' incentive constraints in the auction. The high-type dealer receives a zero payoff against a high-type competitor no matter what her strategy is. If the competitor is uninformed and the customer is a buyer, the high type's payoff from any quote that she provides in period one is at most zero, and thus there are no profitable deviations at this stage. If the competitor is uninformed and the period-one customer is a seller, then the high type's payoff is Δ under any quote strictly higher than $1-\Delta$, whereas any quote lower than that leads to a payoff of at most $\Delta/2$. Thus, quoting 1 at this stage is optimal. In period zero, the high type's overall payoff under an uninformed competitor is $\delta\gamma\Delta/2$. If she deviates to $b<1$ and we set the uninformed competitor's off-path belief to be that this deviation comes from an uninformed type, then the payoff from this deviation is at most $\delta\gamma\min\{\Delta, 1/2\}/2$, which is the continuation payoff from quoting 0 in period one if the future customer is a seller. Thus, this deviation is not profitable. Deviating to $b>1$ leads to the same payoff as quoting her equilibrium strategy.

Similarly, the uninformed dealer never profits from deviating in period one after having played her equilibrium strategy in period zero. Moreover, conditional on the competitor being informed, any quote that she may provide in period zero leads to at most zero payoffs, and thus deviating in period zero against an informed competitor is not optimal. Consider now the uninformed dealer i in period zero when the competitor is uninformed. If dealer i deviates to $b<0$, her payoff in period zero is at most zero. Choose off-path beliefs so that dealer $-i$ believes that dealer i has a low type after observing this deviation. This implies that i's continuation payoff if the future customer is a buyer is zero. If the future customer is a seller, dealer $-i$ will quote $b_u^1 \in (-1, 0)$. If dealer i quotes $b < b_u^1$, she loses the auction and gets zero payoff. If she quotes $b \in [b_u^1, 0]$ and we set the customer's off-path beliefs to zero after observing a pair of quotes (b_u^1, b), then the customer is at best indifferent between trading and not trading following a pair of quotes (b_u^1, b) (recall that she must pay a fee), and thus we can set the customer's strategy to reject trading after this off-path event. Thus, deviating to a lower quote in period zero is not optimal. Deviating to $b \in (0, 1)$ in period zero is also not beneficial if we set the uninformed competitor's belief to assign probability one to the uninformed type after observing this deviation. The analysis of the deviation to $b > 1$ is analogous to the case with $b < 0$.

Finally, we look at the low type. If the competitor is uninformed, dealer i's payoff is Δ in period zero and $\Delta/2$ in period one. This is because, if the period-one customer is a buyer, the uninformed competitor will quote $-1+\Delta$ and, in period zero, the competitor will quote 0. In both cases, the customer will request another quote. Dealer i will then counteroffer $-1+\Delta$, which the customer will accept with probability one. If the future customer is a seller, then the continuation payoff is always zero. If dealer i instead quotes 0, then the competitor will believe that she is uninformed and will quote 0 in period one. If the period-one customer is a buyer, dealer i can then ensure a continuation payoff of $\min\{\Delta, 1/2\}$ if she quotes 0 in period one, or at most Δ if she quotes $-1+\Delta$ and wins the auction with probability one. In period zero, she would get $\min\{\Delta, 1/2\}$. Hence, the payoff from quoting -1 dominates the payoff from quoting 0. Any quote higher than 0 is also suboptimal. Finally, if the competitor has a low type, any quote leads to a zero payoff, and hence the incentive constraint is trivially satisfied.