International Journal of
Financial Studies

MDPI

Article

Anonymity in Dealer-to-Customer Markets

Daniela T. Di Cagno [1], Paola Paiardini [2,*] and Emanuela Sciubba [3]

1 Department of Economics and Finance, LUISS Guido Carli, Viale Romania 32, 00197 Rome, Italy; ddicagno@luiss.it
2 Department of Management, Sapienza University of Rome, Via del Castro Laurenziano 9, 00161 Rome, Italy
3 Birkbeck Business School, University of London, Malet Street, London WC1E 7HX, UK; e.sciubba@bbk.ac.uk
* Correspondence: paola.paiardini@uniroma1.it

Abstract: We use a laboratory experiment to explore the effect of a change in pre-trade anonymity in a quote-driven dealer-to-customer market, organised as a request for quote (RFQ). We consider two treatments in which dealers interact with two types of customers (informed or uninformed). In the first treatment, there is no anonymity: dealers know whether the customer that sent them the request for quote is informed or uninformed. In the second treatment, there is complete anonymity: dealers do not know the type of customers they are interacting with. We find that anonymity improves price efficiency, whereas it does not adversely impact dealers' trading profits. Our results contribute to the debate on transparency versus the adoption of anonymity in financial markets.

Keywords: electronic trading; RFQ; experiments; market microstructure; anonymity

JEL Classification: C90; D40; G10; G14

check for updates

Citation: Di Cagno, Daniela T., Paola Paiardini, and Emanuela Sciubba. 2024. Anonymity in Dealer-to-Customer Markets. *International Journal of Financial Studies* 12: 119. https://doi.org/10.3390/ijfs12040119

Academic Editor: Vassilios Papavassiliou

Received: 15 October 2024
Revised: 22 November 2024
Accepted: 26 November 2024
Published: 29 November 2024

1. Introduction

Many actual financial markets, such as the interbank market, the derivatives market, and the bond market, are organised as multi-dealer electronic markets.[1] Customers simultaneously send a request for a quote (RFQ) to multiple dealers, typically ranging from three to five. The RFQ may disclose the customer's identity, whether they want to buy or sell, the security, quantity, and the settlement date. In response, dealers can provide bids or offers. The customer then selects the best deal, often based on the best price, and the trade is executed shortly after. Rather than engaging in bilateral communication with multiple dealers, an RFQ allows customers to compare prices across dealers easily, promoting competition for client order flow. The RFQ functionality encourages competition among dealers without disseminating trading intentions and dealer quotes to all market participants. In fact, while dealers can see how many others compete, they cannot observe their identity or quotations. This results in better prices while limiting information leakage. RFQ represents the most widely used trading protocol in electronic dealer-to-customer transactions.[2]

One possibility to enhance the resilience of liquidity supply during stressful times is the introduction of all-to-all platforms. This trading mechanism is already establishing itself in the US.

The common belief is that greater transparency will mitigate information asymmetry by improving the price formation process and order execution. Moreover, increased transparency will also provide valuable information to market participants, leading to improved liquidity (Aghanya et al. 2020). Practitioners and academics alike have investigated the potential effects of different transparency regimes both from theoretical and empirical perspectives with mixed results. Technological advancements have transformed the trading protocols of various financial assets. Traditionally organised through bilateral transactions, many assets are now almost exclusively traded electronically, including equities, options, futures, foreign exchanges, and, to a lesser extent, fixed-income securities. Introducing new trading protocols has significantly influenced investor behaviour, improving liquidity, reducing transaction costs, and intensifying provider competition, thereby attracting new market participants.

In this paper, we focus on the impact of disclosing customers' identities and types (whether informed or uninformed) to dealers. Anonymity stands out as perhaps the most controversial aspect of transparency. Especially in the last two decades, there has been a growing tendency to introduce or re-introduce anonymity in financial markets. Most stock exchanges have adopted either pre-trade or post-trade anonymity, or both. Examples include the London Stock Exchange and the Euronext Paris in 2001, the Tokyo Stock Exchange in 2003, the Deutsche Börse AG in 2003, the Italian Stock Exchange (Borsa Italiana) in 2004, the Australian Stock Exchange in 2005, the Helsinki Stock Exchange in 2006, just to name a few. The Italian secondary market for Treasury bonds (MTS) introduced anonymity in 1997.

Considering the heterogeneity of the trading venues in terms of market structure and trading mechanisms, the lack of consensus in the empirical evidence is not surprising. It is therefore essential to assess the implications of changes in the anonymity regime in the context of a specific trading environment, which makes it possible to identify the actual mechanisms through which anonymity (or the lack thereof) affects market quality.

A non-anonymous market setting may enable traders to enforce informal collusive agreements to quote wider spreads.[3] Alternatively (or in addition), when market participants are asymmetrically informed, anonymity may affect market quality because of the differential benefits that it brings to different types of traders. If presented with a choice of platforms, informed dealers might migrate to anonymous trading venues (Reiss et al. 2005), thereby increasing their trading volume and liquidity. Furthermore, in a non-anonymous setting, informed traders may dissimulate their information by setting a wider spread, in an attempt to reduce the free-riding from uninformed traders who would otherwise infer information from market prices, thereby reducing informed traders' advantage (Foucalt et al. 2007). Finally, anonymity may increase the incentives to acquire information when information acquisition is endogenous, thereby augmenting the number of informed traders, and ultimately increasing liquidity and reducing the spread (Rindi 2008). In a very recent experimental paper, Ruiz-Buforn et al. (2021) analyse the impact of public disclosures in the presence of endogenous private information. They find that releasing public information does not necessarily improve market efficiency when private information is costly, and that the way public information is released must be carefully considered since it has relevant consequences for market performance and price efficiency.

Limited experimental research exists on pre-trade transparency. Gozluklu (2016) explores a market with hidden orders, finding that anonymity enhances liquidity by helping traders reduce order exposure. Our study aligns with these findings, focusing on an alternative mechanism through which pre-trade transparency may impede market quality, particularly the adverse selection problem faced by dealers when trading with informed customers. In contrast, Perotti and Rindi (2006) investigates the impact of anonymity in an electronic open-book market. They find that disclosing traders' identities reduces the incentive to acquire information, leading to fewer informed traders and a subsequent decrease in market liquidity. On the other hand, Majois (2008) considers an order-driven market with a fixed share of informed traders and concludes that market anonymity does not affect liquidity and efficiency.

The variation in anonymity regimes on real trading platforms serves as a natural experiment, enabling an assessment of transparency changes within the same platform. However, the exact impact of anonymity on market outcomes, such as lower liquidity and increased spreads, remains unclear due to potential factors like collusive agreements and asymmetric information. The theoretical, empirical, and experimental literature on transparency in securities markets primarily focuses on limit-order markets.

The gap in the literature that this paper aims to address involves employing experimental markets to investigate the impact of anonymity in a multi-dealer RFQ trading system.[4] Dealers play a crucial role in over-the-counter (OTC) markets, often owning the platforms where they trade with customers. We argue that introducing anonymity could increase customer participation in these platforms, redirecting orders from OTC to electronic trading. This shift may benefit small customers by providing better quotes typically reserved for larger customers and facilitating the execution of large trades without compromising dealers' profits, potentially avoiding winner's curse scenarios. To the best of our knowledge, this is the first study on anonymity in an RFQ market, and this is our main contribution to the existing literature.

In our experimental setup, customers seek quotes from dealers for trading a single asset. Unlike prior studies, we disregard the decision on information acquisition and focus solely on how anonymity influences market quality with a constant number of informed traders. Two customer types exist, as follows: informed customers know the asset's true value, while uninformed customers do not. Unaware of the true value, dealers compete in a first-price auction when responding to customer requests. We evaluate two treatments, namely the Transparent market, where dealers observe customer types before quoting, and the Opaque market, where they cannot discern the customer type.

The advantage of the experimental setting lies in the ability to control factors that would be uncontrollable in a real scenario. However, it is important to note that we can only manipulate one variable at a time. Otherwise, we cannot confidently attribute the observed effect to the change in that specific variable. In our study, we refrain from altering the trading setting of the RFQ, except for the introduction of anonymity. While this approach has its limitations, particularly in the context of all-to-all platforms where dealer competition is higher and access to other dealers' quotes is possible, we contend that the scenario remains relatively unchanged from the customers' perspective.

Our findings show that anonymity enhances price efficiency without negatively impacting dealers' profits when comparing the Opaque and Transparent platforms. The increased adoption of multi-dealer platforms could facilitate clients in comparing multiple price sources, potentially fostering greater dealer competition, which may, in turn, influence liquidity provision.

The reminder of this paper is organised as follows. Section 2 explains the organisation of the dealer-to-customer electronic trading. Section 3 describes the experimental design and procedures. Section 4 discusses the experimental results. Section 5 concludes the paper.

2. Dealer-to-Customer Electronic Trading

Technology transformed the trading process for a variety of financial assets. Many markets that were organised through bilateral transactions are now almost exclusively electronic, such as equities, options, futures, and foreign exchanges; for some other markets, such as the fixed-income market, the percentage of transactions taking place in the over-the-counter (OTC) market is still notable. Nevertheless, electronic trading has become an increasingly important part of the fixed-income market as well in recent years and, for some fixed-income securities, 'electronification' has reached a level similar to that of other markets. Electronic trading is having a significant impact on the structure of the markets, leading to measurable improvements in transaction costs, price discovery, liquidity measures, and other market-quality metrics. Moreover, these venues have allowed the participations of new market participants. The introduction of new trading protocols has influenced investors' behaviour and liquidity, and increased competition among providers.

Trades not realised via bilateral negotiations occur on electronic platforms. These platforms are called alternative trading systems (ATSs) in the U.S. and Canada, and a multilateral trading facility (MTF) in Europe. Examples of trading platforms organised in this way include, among others, Tradeweb, the global operator of electronic marketplaces for rates, credit, equities, and money markets; BondVision for European bond transactions; CanDeal for Canadian government bonds; MarketAxess for global corporate bonds; and BrokerTec Quote for the European repo market.

Using our experimental framework, we aim to investigate the effect of introducing anonymity in a multi-dealer-to-customer platform. Details on the trading structure are illustrated in Figure A1 in Appendix A. The experimental setting is illustrated in the next section.

3. Experimental Design and Procedures

3.1. Experiment Overview

The experimental environment is organised as a multi-dealer request for quote (RFQ) market with twenty dealers and four customers. Each customer sends a simultaneous request to trade one unit of the asset (either to buy or to sell) to five different dealers. Experimental subjects always play the role of dealers, while customers are computers.

The market is divided into twelve trading rounds of thirty seconds each. At the beginning of each trading round, dealers are provided with 150 Experimental Currency Units (ECUs) and one asset unit, and compete with the other four participants in the experimental session to either sell or buy one asset unit.

We use a common value design, i.e., the true value of the asset is identical for all subjects and does not change from one round to another. At the beginning of the experiment, participants are informed that the true value of the asset is uniformly distributed between a minimum of 50 ECUs and a maximum of 150 ECUs. Given these bounds, a uniform prior expectation of the true value would be approximately 100 ECUs. Thereafter, participants can revise their prior through the trading rounds, when they are able to trade. In line with the existing literature (Bloomfield and O'Hara 1999; Flood et al. 1999), we consider different true values, distinguishing between prices both far away and close to the uniform prior expectation of 100 ECUs[5].

We implement a between-subjects experimental design using two different transparency settings: Opaque and Transparent. In the Opaque treatment conditions, dealers do not know whether the customers with whom they are interacting are informed, i.e., know the true value of the asset, or uninformed, i.e., they do not know it. In the Transparent treatment conditions, dealers are instead informed about the type of customers (informed or uninformed) they are dealing with.

The experiment was programmed using z-Tree (Fischbacher 2007) and was conducted at the Behavioural, Experimental and Data Science Laboratory (BEADS Lab), University of Birmingham. We used the ORSEE (Greiner 2004) enrolment platform to recruit 160 participants among university undergraduate students. All subjects participated in the experiment for the first time and were immediately paid at the end of each session. Sessions lasted an average of 45 min and the average payment was GBP 7. Written instructions were handed to participants and read aloud before the beginning of each session of the experiment (see Appendix B).

Table A1 in Appendix A reports the balance checks of subjects' characteristics, indicating no differences in the distribution of sociodemographic variables between treatments.

3.2. The Trading Mechanism

Dealers are randomly divided into groups of five. We adopt a stranger-matching protocol so that new groups of five are randomly formed at the beginning of each round.

All the dealers in a group receive the same request for a quote, either to buy or to sell, from a customer (that could be either informed or uninformed according to the treatment), and they have to respond with a price proposal. When they send their price proposal, they are aware that they are in competition with four other dealers, of whom they do not know either the identity (the type) or the proposed prices. Therefore, the underlying trading mechanism is a type of first-price sealed bid auction: the dealer that makes the highest (lowest) price proposal wins the auction and buys (sells) the asset.

In both treatments, dealers know that there is the same proportion of informed and uninformed customers and that both types of customers will buy or sell with equal probability. Moreover, they are informed about the trading strategies implemented by the two types of customers: uninformed customers will always buy at the lowest available price, and sell at the highest available price; informed customers will buy only if the lowest available price is strictly lower than the true value of the asset, and they will sell only if the highest available price is strictly higher than the true value of the asset. In our design, dealers are not able to opt out of trading by not quoting; however, they are able to avoid trading by making extreme quotes which will not be accepted by informed customers and therefore result in no exchange for that round.

At the end of every trading round, each dealer observes whether s(he) was able to buy (sell) the asset and the highest (lowest) price at which it has been sold (bought).

3.3. Hypotheses and Expected Results

We compare Transparent and Opaque markets along three basic dimensions: trading frequency, price efficiency, and dealers' profits.

Dealers always lose by trading with an informed customer, and we would therefore expect no trade to take place with informed customers in the Transparent treatment, for a mechanism akin to the classic 'no trade' results in Grossman and Stiglitz (1980) and Milgrom and Stokey (1982). In our design, dealers are not able to opt out of trading by not quoting; however, they are able to avoid trading by making extreme quotes which will not be accepted by informed customers and therefore result in no exchange for that round.

Hypothesis 1. *The trading frequency with informed customers in the Transparent treatment is lower than in the Opaque treatment.*

When dealers trade less often with informed customers, they have fewer opportunities for learning the true value of the asset.

Hypothesis 2. *Price efficiency is greater in the Opaque treatment than in the Transparent treatment.*

In the Transparent treatment, dealers are more able to shield themselves from losses by avoiding trades with informed customers; at the same time, they learn less, and this may affect their profits.

Hypothesis 3. *There is no significant difference between dealers' profits across the two treatments.*

4. Empirical Results

In this section, we illustrate our experimental results. We first compare the two treatments by analysing the average number of trades.

We will then proceed by implementing a between-subjects repeated-measures ANOVA design, which is a well-established procedure for analysing experimental data, widely used in the literature (Bloomfield et al. 2005; Chelley-Steeley et al. 2015), to analyse price efficiency and trading profits. For each dependent variable that we take into consideration, we examine two between-subject components and one repeated component. The between-subject components, which are fixed, are: treatment (two categories: Opaque and Transparent) and true value (four categories: 68, 86, 109, and 142). The repeated component is given by the timing (three categories: early (1–4 rounds), middle (5–8 rounds) and late (9–12 rounds)). We also analyse the interaction effects between the independent variables; this allows us to observe whether the effects of timing and true value differ between treatments. Finally, we consider one-way ANOVA to analyse the effect of the timing distinguished by true value for each treatment.

4.1. Trading Frequency

In our experimental setting, dealers can trade only one unit of the asset. The average trading frequency, measured as the number of times a transaction is concluded, is higher in the Transparent treatment (2.30) than in the Opaque treatment (2.14). A Mann–Whitney test shows that this difference is statistically significant ($p = 0.035$)[6]. Moreover, we further analyse the trading frequency, distinguishing between the type of customer (informed or uninformed) that the dealers interact with, for each treatment. We find that the average trading frequency in the Transparent treatment is 2.50, when dealers conclude a transaction with an uninformed customer, which is significantly higher (p = 0.008) than the trading frequency when they conclude a transaction with an informed one, at 2.09. This is an expected result, as dealers typically avoid trading with informed customers who know the true value of the asset. In line with the well-known 'no trade' results (Grossman and Stiglitz 1980; Milgrom and Stokey 1982), in the Opaque treatment, the average trading frequency when dealers interact with an uninformed customer is 2.17, and, when they interact with an informed one, it is 2.09. The difference is not statistically significant (p = 0.193), as we would have expected, given that the type of customer is unknown to dealers in this setting. These average trading frequencies are illustrated in Figure 1.

We can conclude that the trading frequency is higher in the Transparent treatment, where dealers can distinguish between the type of customers, with respect to the Opaque one. Moreover, within the Transparent treatment, trading frequency is higher when dealers interact with uninformed customers. In relative terms, the frequency of trades with informed customers is lower in the Transparent treatment than in the Opaque, which confirms Hypothesis 1.

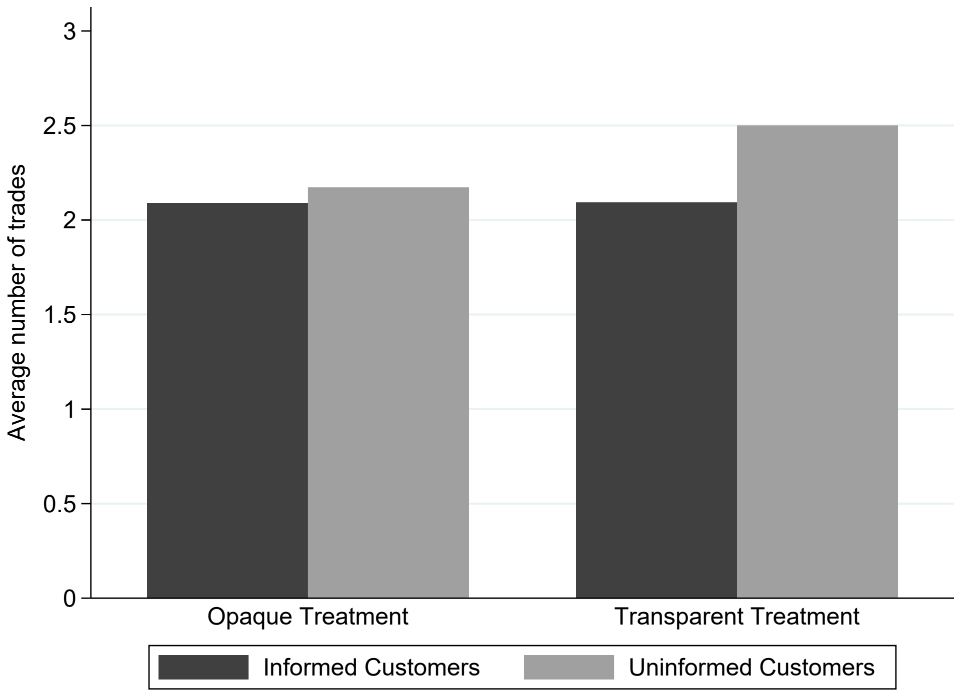

Figure 1. Trading frequency by treatment and customer type. The graph illustrates the average number of trades for the Opaque and Transparent treatment made by informed and uninformed customers.

4.2. Price Efficiency

We define price efficiency as the absolute deviation between the price and the asset's true value divided by the asset's true value to make the different true values comparable.

$$\text{Price Efficiency} = \sum_{k=1}^{K} \sum_{t=1}^{T} \frac{|P_{kt} - TV_k|}{TV_k} \tag{1}$$

where P_{kt} is the winning proposal for the market k during round t and TV_k is the true value for session k, which remains the same throughout the rounds.

Figure 2a shows the price efficiency across rounds for the two treatments. We can observe that in both treatments, the average pricing error remains quite constant up to round six, and then we detect a decline, although this is more pronounced for the Opaque treatment. Finally, for the last three rounds, we observe a more irregular pattern for both treatments. These findings suggest that price efficiency improves over time, resulting in more accurate pricing as participants gain experience throughout the rounds. Figure 2b shows that the evolution of average pricing error through rounds for informed customers closely resembles the pattern observed for uninformed customers.

Panel A of Table 1 reports mean price efficiency data per true value and timing, averaged by treatment. We can observe that the average price efficiency is better for less extreme asset values (i.e., closer to the expected value of 100) and for later rounds. This is because discovering the true value is easier when the value is closer to the expected value, and in later rounds, when participants receive feedback after the trade. Panel B of Table 1 presents the ANOVA results. The average pricing error in the Transparent treatment (0.34) is higher than in the Opaque one (0.29) and their difference is statistically significant ($p < 0.001$)[7], indicating that the information setting plays a role for price efficiency, which is better in the less transparent venue. This result confirms Hypothesis 2.

When we focus our attention on each treatment separately, we find that both timing and true value influence price efficiency. In particular, true value is significant in both treatments (p 0.003), irrespective of the time, and it decreases as the time passes in both treatments ($p < 0.001$). Moreover, the interaction between timing and true value is significant for both treatments ($p = 0.029$), so we are going to analyse how price efficiency changes, through the passage of time, for more and less extreme values. In the Opaque treatment for more extreme values and in the Transparent treatment for less extreme values, there is a significant difference between the early and late time intervals.

Price efficiency for each session is reported in Figure A2 in the Appendix A.

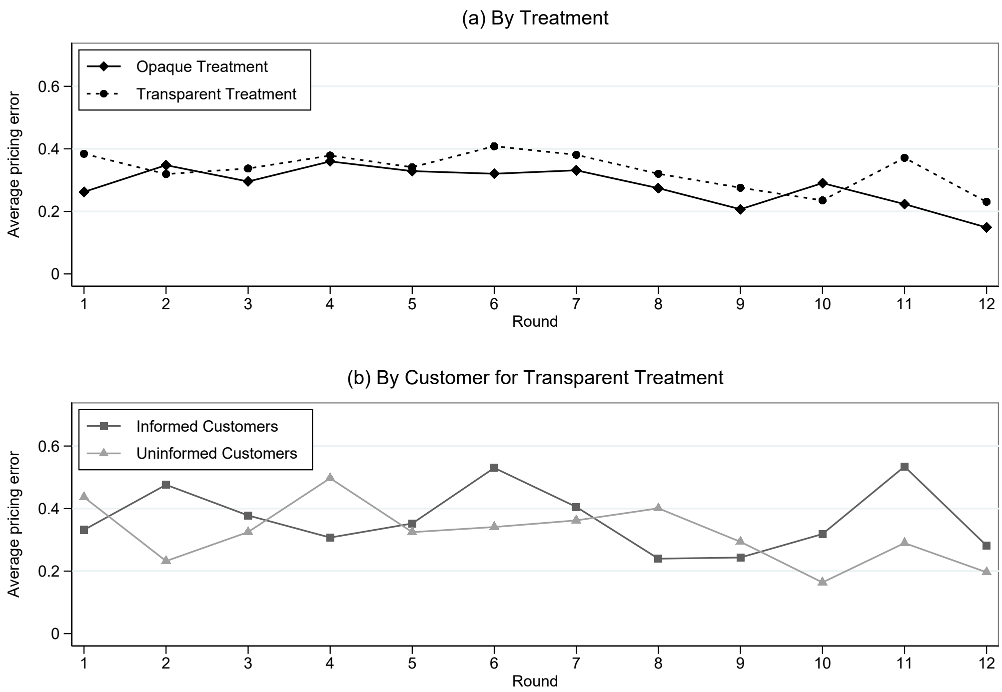

Figure 2. Price efficiency by treatment and customer type. Graph (**a**) illustrates plots of average pricing error in the two treatments for each round. Graph (**b**) displays the average pricing error for each customer in the Transparent treatment, for each round.

Turning our attention to the ANOVA by customer for the Transparent treatment, we find that price efficiency significantly differs between the two types of customers ($p < 0.001$) and it is better when dealers meet uninformed customers. The true value is significant for both customers ($p < 0.001$), irrespective of the time. For uninformed customers, there is a significant difference in price efficiency both between middle and late and early and late time intervals. This means that price efficiency improves in the subsequent phases with respect to the early interval (Table 2).

Table 1. Price efficiency.

Panel A. Means by true value and timing

		By True Value				By Timing		
	Average	69	86	109	142	Early	Middle	Late
Opaque	0.29	0.37	0.21	0.23	0.35	0.32	0.32	0.23
Transparent	0.34	0.43	0.34	0.21	0.38	0.36	0.36	0.29

Panel B. ANOVA

All True Values

	Opaque Treatment				Transparent Treatment			
		By Timing				By Timing		
	All	Early	Middle	Late	All	Early	Middle	Late
True Value	<0.001	<0.001	<0.001	0.003	<0.001	<0.001	<0.001	<0.001
Timing	<0.001				<0.001			
Timing × True Value	<0.001				0.029			

By True Value: Timing Main Effect

	Opaque Treatment				Transparent Treatment			
	All	Early vs. Middle	Middle vs. Late	Early vs. Late	All	Early vs. Middle	Middle vs. Late	Early vs. Late
Less Extreme Values	0.128	0.314	0.103	0.823	0.001	0.941	0.002	0.001
More Extreme Values	<0.001	0.201	<0.001	<0.001	0.157	0.900	0.117	0.268

Panel A reports average price efficiency per true value and timing, by treatment. Panel B reports the ANOVA results for the effect of true value and timing in the All column, separately for each treatment. The Early, Middle, and Late columns report separate ANOVA results using only the true value for that subperiod, respectively. Two-tailed pair-wise tests of timing differences for each treatment are given in the columns labelled Early vs. Middle, Middle vs. Late, and Early vs. Late.

Table 2. Price efficiency—Transparent treatment.

Panel A. Means by true value and timing

		By True Value				By Timing		
	Average	69	86	109	142	Early	Middle	Late
Informed Customer	0.36	0.38	0.37	0.21	0.49	0.36	0.36	0.34
Uninformed Customer	0.32	0.45	0.30	0.21	0.31	0.35	0.36	0.24

Panel B. ANOVA

All True Values

	Informed Customer				Uninformed Customer			
		By Timing				By Timing		
	All	Early	Middle	Late	All	Early	Middle	Late
True Value	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Timing	<0.001				<0.001			
Timing × True Value	0.011				<0.001			

By True Value: Timing Main Effect

	Informed Customer				Uninformed Customer			
	All	Early vs. Middle	Middle vs. Late	Early vs. Late	All	Early vs. Middle	Middle vs. Late	Early vs. Late
Less Extreme Values	0.022	0.205	0.026	0.645	0.007	0.767	0.019	0.003
More Extreme Values	0.001	0.005	0.030	0.926	<0.001	0.011	0.013	<0.001

Panel A reports average price efficiency per true value and timing, by customer. Panel B reports the ANOVA results for the effect of true value and timing in the All column, separately for each customer. The Early, Middle, and Late columns report separate ANOVA results using only the true value for that subperiod, respectively. Two-tailed pair-wise tests of timing differences for each customer are given in the columns labelled Early vs. Middle, Middle vs. Late, and Early vs. Late.

4.3. Trading Profits

Figure 3a shows the average payoff across rounds for the two treatments. On average, dealers in both treatments experience losses and there is not a statistically significant difference between the two treatments ($p = 0.933$). Figure 3b illustrates the average payoff through rounds for each customer in the Transparent treatment. Also, in this case, we do not inspect significant differences. These depictions are confirmed by the ANOVA results reported in Table 3.

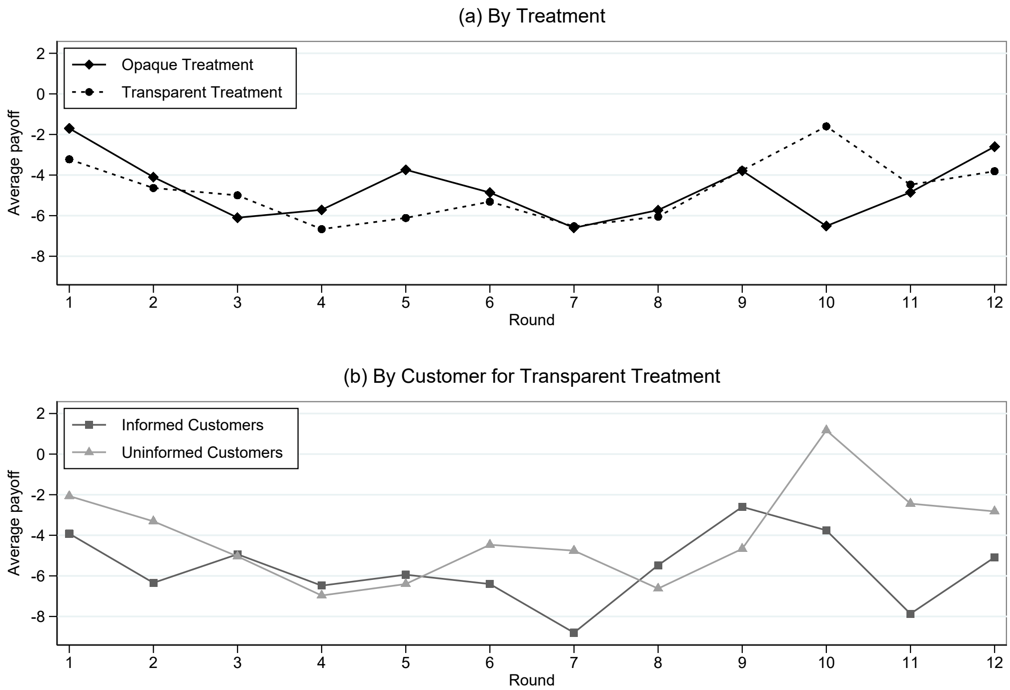

Figure 3. Payoff by treatment and customer type. Graph (**a**) illustrates plots of average dealer's payoff in the two treatments, for each round. Graph (**b**) displays the average dealer's payoff for each customer in the Transparent treatment, for each round.

Panel (A) of Table 3 presents the mean payoff for true value and timing. Panel (B) shows that there are no differences in payoff for different true values and for different timing in both treatments. We can conclude the same for customers' payoff, which is the opposite of dealers' payoff. We also repeated the ANOVA analyses, distinguishing the type of customer in the Transparent treatment; however, we have not reported the results, since we did not find any significant differences. These results are in line with Hypothesis 3.

Table 3. Dealers'payoff.

Panel A. Means by true value and timing								
		By True Value				By Timing		
	Average	69	86	109	142	Early	Middle	Late
Opaque	−4.69	−4.42	−5.73	−3.97	−4.64	−4.40	−5.23	−4.44
Transparent	−4.76	−4.70	−5.44	−4.55	−4.38	−4.88	−6.00	−3.41
Panel B. ANOVA								
All True Values								
	Opaque Treatment				Transparent Treatment			
		By Timing				By Timing		
	All	Early	Middle	Late	All	Early	Middle	Late
True Value	0.712	0.496	0.943	0.351	0.950	0.827	0.993	0.590
Timing	0.755				0.106			
Timing ×True Value	0.571				0.795			

Table 3. *Cont.*

By True Value: Timing Main Effect								
	Opaque Treatment				Transparent Treatment			
	All	Early vs. Middle	Middle vs. Late	Early vs. Late	All	Early vs. Middle	Middle vs. Late	Early vs. Late
Less Extreme Values	0.775	0.819	0.998	0.785	0.539	0.964	0.455	0.615
More Extreme Values	0.638	0.939	0.603	0.807	0.275	0.549	0.166	0.720

Panel A reports average payoff per true value and timing, by treatment. Panel B reports the ANOVA results for the effect of true value and timing in the All column, separately for each treatment. The Early, Middle, and Late columns report separate ANOVA results using only the true value for that subperiod, respectively. Two-tailed pair-wise tests of timing differences for each treatment are given in the columns labelled Early vs. Middle, Middle vs. Late, and Early vs. Late.

5. Conclusions

This paper uses an experimental approach to assess how changing the level of pre-trade anonymity impacts important market features, such as trading frequency, price efficiency, and trading profits in a dealer-to-customer request for quote market.

In the Transparent treatment, subjects know the type of customers, whether informed and uninformed. In the Opaque treatment, trades are anonymous, i.e., subjects do not know the type of customers they are matched with. Dealers know, when they send a price proposal, that they are in competition with other four dealers in their group, from whom they can observe neither the type nor the proposed prices. Our main findings indicate that the Opaque treatment, i.e., the one with pre-trade anonymity, improves trading frequency with informed customers and price efficiency, without negatively impacting dealers' profits.

The experimental method allows us to analyse the dealers' strategies in the two treatments, without other behavioural confounders. Firstly, we do not take into consideration the fact that dealers have usually access to multiple venues, so that they can implement hedging strategies. Secondly, our experimental approach allows us to ignore any recursive client relationships with customers, which might influence dealers' quoting behaviour.

The experimental approach does not capture in full the complexity of real-world financial markets, but it allows us to contribute to the debate on transparency in financial markets by isolating the effect of alternative anonymity regimes on dealers' behaviour. To the best of our knowledge, this paper represents the first experimental study on dealer-to-customer markets, filling a gap in the literature but also providing insights into the effects of possible policy changes in these markets. In light of our findings, we offer interesting insights on how introducing anonymity in the dealer-to-customer settings might increase the participation of small customers, usually trading over the counter, because they are discouraged by the competition with more important clients in a transparent setting. Overall, our study highlights that the concern that informed dealers would be adversely affected by anonymity, deciding to opt out of the market, is unfounded and prompts for further research in this under-investigated area through larger-scale field experiments.

Author Contributions: Conceptualization, D.T.D.C., P.P. and E.S.; methodology, P.P.; software, P.P.; validation, D.T.D.C., P.P. and E.S.; formal analysis, P.P.; investigation, P.P.; data curation, P.P.; writing—original draft preparation, D.T.D.C., P.P. and E.S.; writing—review and editing, D.T.D.C., P.P. and E.S.; visualization, P.P.; funding acquisition, D.T.D.C. and P.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the University of Birmingham and LUISS Guido Carli.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: The experiment was approved by the Birmingham Research Ethics Committee, and participants signed an informed consent form before taking part in the experiment.

Data Availability Statement: Data are available from the authors upon request.

Conflicts of Interest: The authors have no competing interests to declare that are relevant to the content of this article.

Abbreviations

The following abbreviations are used in this manuscript:

RFQ	Request for quote
OTC	Over-the-counter
ATS	Alternative trading systems
MTF	Multilateral trading facility
SDPs	Single-dealer platforms
MDPs	Multi-dealer platforms

Appendix A

The request for quote (RFQ) protocol is a long-established execution mechanism in the dealer-to-customer segment that allows customers to request quotes from dealers. It can take two forms: single-dealer platforms (SDPs) and multi-dealer platforms (MDPs). The former are proprietary electronic trading systems offered by individual dealers for trading with their customers. The latter allow for a customer's request for a quote to be sent simultaneously and instantly to multiple dealers (usually between three and five) instead of just one. The request reveals to the dealers the name of the customer, whether it is a buy or sell, the security, the quantity, and the settlement date; dealers, in turn, can respond with bids or offers. The customer chooses the deal that they like best (typically the best price), if any, and the trade is executed shortly after. Running an RFQ auction, rather than contacting multiple dealers bilaterally, enables the customer to easily compare prices across dealers and thus promotes competition for client order flow among dealers. RFQ functionality encourages dealers' competition without disseminating trading intentions and dealer quotes to all market participants, since dealers can see how many other dealers compete, but they cannot observe either their identity or their quotations. This results in better prices while limiting information leakages.

RFQ protocols may differ in one or more of the following aspects: whether the quote requester or receiver reveals their identity; whether the type of order (buy or sell) is disclosed; how many and what type of participants may receive the RFQs; and whether the quotes are executable or indicative. Many real-world financial markets, such as the interbank, derivatives, and bond markets, are organized as multi-dealer electronic markets.

The most common type of trading organization consists of two distinct segments: a multi-dealer-to-customer segment and an inter-dealer segment. In the dealer-to-customer segment, customers are matched with dealers but not with other customers. Dealers trade with one another on a platform organized as an anonymous limit-order book, which is inaccessible to customers.

Dealers have a strong influence on the OTC markets and often they own the platform where they trade with customers. In particular, we argue that introducing anonymity could further increase the participation of customers on these platforms, diverting orders from OTC to electronic trading, in at least in two ways: by allowing small customers to receive better quotes that are usually reserved to bigger customers; and by allowing for the execution of large trades (which could instead generate winner's curse[8]), without reducing dealers' profit.

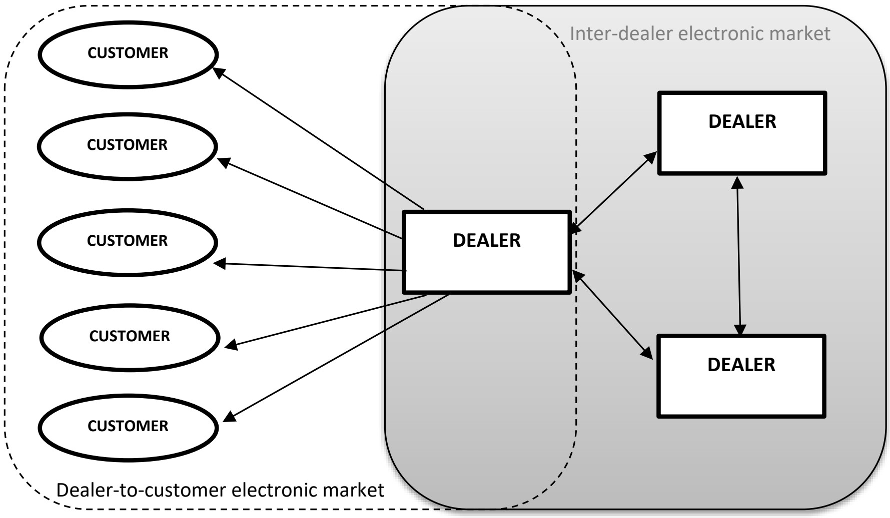

Figure A1. Electronic trading structure.

Table A1. Balance checks of individuals' characteristics across treatments.

	Opaque Treatment	Transparent Treatment	Mann–Whitney Test
Female	0.39	0.50	0.153
	(0.49)	(0.50)	
Age	23.85	22.48	0.568
	(5.77)	(3.44)	
Married	0.19	0.11	0.185
	(0.39)	(0.32)	
Religion	0.53	0.55	0.752
	(0.50)	(0.50)	
Economics	0.44	0.40	0.632
	(0.50)	(0.49)	
Nationality	0.36	0.34	0.741
	(0.48)	(0.47)	
Trust	0.30	0.24	0.374
	(0.46)	(0.43)	
Risk	0.54	0.46	0.429
	(0.50)	(0.02)	

Note: The table reports the average values of each demographic characteristic, along with the corresponding standard deviations for each treatment: Female, Age, Married, Religion, Economics, Nationality, Trust, and Risk. Female is a dummy variable, taking the value 1 if a subject is female and 0 otherwise. Age corresponds to the average subjects' age. Married is a dummy variable, equal to 1 if the respondent is married and 0 otherwise. Religion is a dummy variable, taking the value 1 if the respondent declares belonging to some kind of religion and 0 otherwise. Economics is a dummy variable, equal to 1 if subjects study an economics degree and 0 otherwise. Nationality is equal to 1 for UK nationality and 0 otherwise. Trust is a dummy variable, taking the value 1 if subjects agree to the statement that people can be trusted and 0 otherwise. Risk is a self-assessment measure of attitude toward risk. The responses are evaluated on an eleven-point Likert scale ranging from 0 ("I love risk-taking behaviour") to 10 ("I am a person always avoiding risk"). We transform it into a dummy variable, which is equal to 1 for values up to 5 on the Likert scale, and 0 otherwise. The last column of the table reports the p-values from two-sided Mann–Whitney tests, indicating that subjects' observable characteristics were similar between treatments.

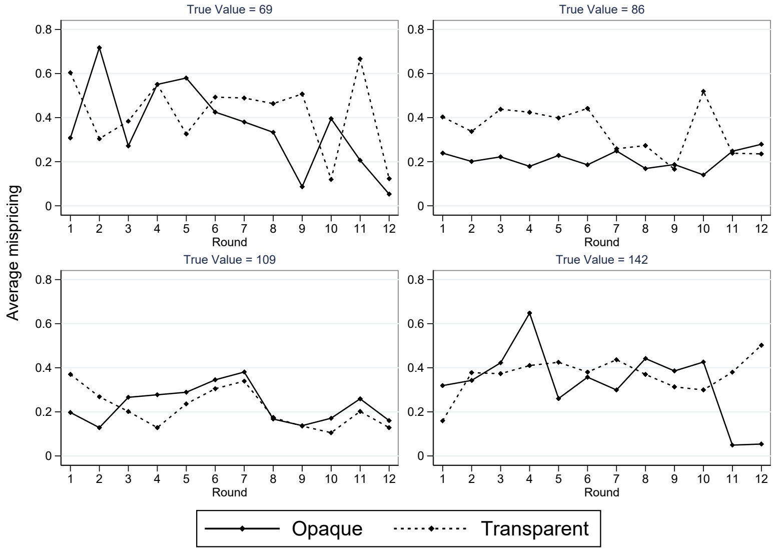

Figure A2. Price efficiency by treatment and true value.

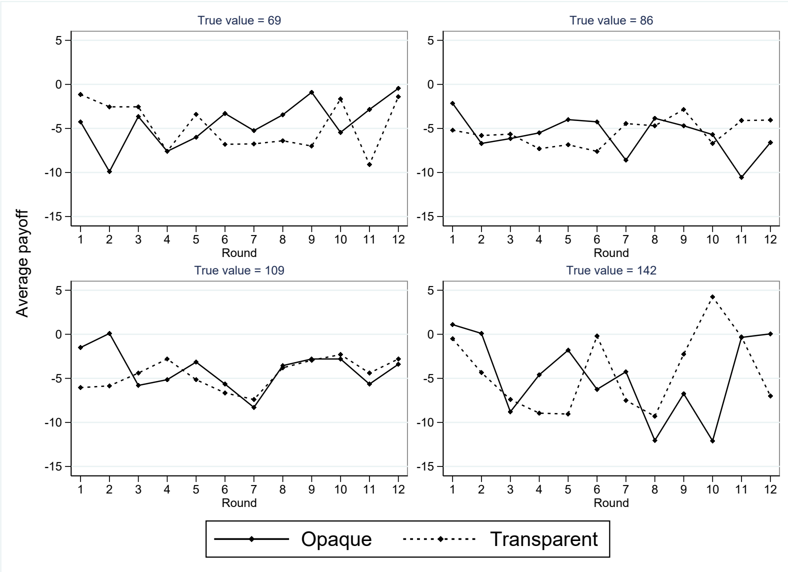

Figure A3. Payoff by treatment and true value.

Appendix B. Experiment Instructions

Introduction

Welcome! You are about to take part in a decision-making experiment. This experiment is run by the 'Birmingham Experimental Economics Laboratory' and has been funded by various research foundations. Just for showing up you have already earned £5.00. You can earn additional money depending on the decisions made by you and other participants. It is therefore very important that you read these instructions with care.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your following of these rules.

We will first jointly go over the instructions. After we have read the instructions, you will have time to ask clarifying questions. Then you will fill in a short questionnaire to check your understanding of the experiment.

We would like to stress that any choices you make in this experiment are entirely anonymous. Please do not touch the computer or its mouse until you are instructed to do so. Thank you.
In the instructions, unless otherwise stated, we will not speak in terms of British Pounds, but in terms of Experimental Currency Units (ECUs). Your entire earnings will, thus, be calculated in ECUs. At the end of the experiment the total amount of ECUs you have earned will then be converted into British Pounds at the following rate: 1ECU= £0.10.
The converted amount will be privately paid to you in cash at the end of the experiment.
Detailed Information about the Experiment
You will operate on a market where you will trade an asset. At the beginning of the experiment the computer will randomly extract the true value of the asset. This value is an integer number between 50 and 150 ECUs, included. Each value is equally likely. The value will remain the same throughout the experiment.

Your Role in the Experiment

In this market there are 20 Dealers and 4 Customers. Dealers will never be informed of the true value of the asset. Customers can be of two types: 50% of the customers are informed, i.e. they know the true value of the asset, and 50% are uninformed, i.e. they do not know the true value of the asset. Customers are always computers. Each of you will act as a Dealer.

Trading Rules

The market is composed of 12 rounds. In each round you will be endowed with one asset and 150 ECUs. Asset and money at the end of a round are not carried forward to the next round.
At the beginning of each round, you will be randomly matched in groups of five. The group composition will change from one round to another. At no point during the experiment, nor afterwards will you be informed about the identity of the other participants in your group, and the other participants will never be informed about your identity.
All the dealers in a group will receive the same request for quote from a customer, either to buy or to sell, and will have to respond with a price quotation. You will not know whether the customer who sent you the request is informed or uninformed. Moreover, you will never be informed of the price quotations of the other dealers in your group. You will enter your decisions on a screen like the one shown below.

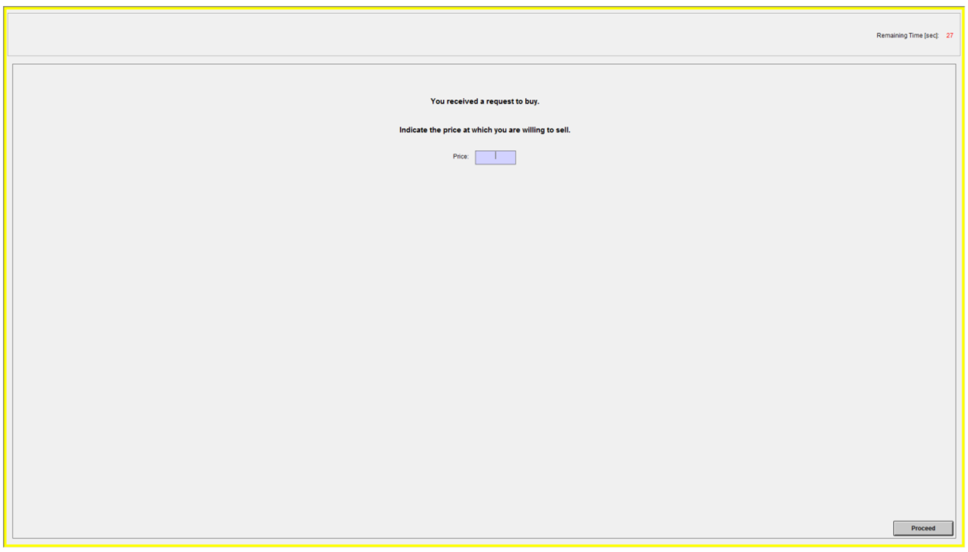

During each round, you are continuously informed of the remaining time until the end of the round (at the top right of the screen-shot: ’remaining time’). Customers will buy or sell with a probability of 50%. They will see all the quotes of the five dealers and will choose the best one, according to their trading rule.

Uninformed customers adopt the following trading rule: they will always buy at the lowest available price and they will always sell at the highest available price.

Informed customers adopt the following trading rule: they will buy only if the lowest available price is strictly lower than the true value of the asset, and they will sell only if the highest available price is strictly higher than the true value of the asset.

Thus, it might be possible that during a round there are no exchanges.

At the end of a trading round you will see whether you were able to buy or sell the asset, the highest price at which it has been sold or the lowest price at which it has been bought.

Making Money

Earnings at the end of each round are determined by the difference between the price at which you were able to trade the asset, and the true value. Every asset bought will result in a gain or a loss given by: (True Value-Price). Every asset sold will result in a gain or a loss given by: (Price-True Value).

Example: assuming that the true value of the asset was 70. In a given round, buying at a price of 65 will result in a gain of $(70 - 65) = 5$, selling at a price of 65 results in a loss of $(65 - 70) = -5$.

Price Estimates

At the beginning of each round, you have to input a price estimate of the true value of the asset. Please pay attention to indicate your estimate since, at the end of the experiment, only one round will be randomly extracted and, those who for that round indicated a price estimate within $+/-10\%$ of the true value will receive £3.00.

Your total earnings at the end of the experiment are computed as follows:

Show-up fee + Earnings from the market game + Earnings from the price estimate.

Do you have any questions? Please raise your hand and an experimenter will come to your desk. Please do not ask any question out loud.

Notes

1 Examples of trading platforms organised in this way include, among others, Tradeweb, the global operator of electronic marketplaces for rates, credit, equities, and money markets; BondVision for European bond transactions; CanDeal for Canadian government bonds; MarketAxess for global corporate bonds; and BrokerTec Quote for European repo market.

2 According to a survey on electronic trading platforms, dealer-to-customer platforms accounted for 95% of the total transaction volume in 2014 (BIS 2016).

3 Barclay et al. (2003); Simaan et al. (2003) study the competition between anonymous Electronic Communication Networks (ECNs) and non-anonymous Nasdaq dealers. Barclay et al. (2003) find that ECNs attract informed traders for Nasdaq listed stocks as they offer speed of execution, and pre- and post-trade anonymity. (Simaan et al. 2003) find that competition is stronger and quotes are tighter on the anonymous ECNs.

4 O'Hara and Zhou (2021) analyse how trading corporate bonds on MarketAxess, an electronic request-for-quote platform, impacts dealers and trading dynamics. While dealers remain pivotal, the study suggests the potential for further shifts toward all-to-all trading platforms. Similarly, Allen et al. (2023) examine the regulatory question of centralising over-the-counter (OTC) markets to centralised platforms in the Canadian government bond market. Their model predicts that overall welfare would increase despite a persistent price gap between large and small investors on centralised platforms.

5 We randomly extract four prices from each of the following four intervals: [50, 75), [75–100), [100–125), and [125, 150]. The four random true values are: 69, 86, 109, and 142.

6 We always refer to a two-tail test, unless otherwise specified.

7 This result is obtained by combining the two treatments into a single ANOVA, incorporating a categorical treatment factor.

8 In the dealer-to-customer segment, the selected dealers submit a quote proposal in response to a client request. The winning dealer can access the inter-dealer market and hedge their risk. However, the dealers who don't win can also assume a contrary position in the inter-dealer segment. Thus, the winning dealer does not necessarily have an advantage with respect to the others.

When large trades are executed and there is a delay between the request and execution of the trade, this problem is exacerbated. See Dunne et al. (2006) for a discussion related to the trading organisation of the European government bond market.

References

Aghanya, Daniel, Vineet Agarwal, and Sunil Poshakwale. 2020. Market in Financial Instruments Directive (MiFID), stock price informativeness and liquidity. *Journal of Banking & Finance* 113: 105730.

Allen, Jason, and Milena Wittwer. 2023. EconPapers: Centralizing Over-the-Counter Markets? *Journal of Political Economy* 131: 3310–51. [CrossRef]

Barclay, Michael J., Terrence Hendershott, and D. Timothy McCormick. 2003. Competition among trading venues: Information and trading on electronic communications networks. *The Journal of Finance* 58: 2637–65. [CrossRef]

Bank for International Settlements. 2016. Electronic Trading in Fixed Income Markets. BIS Working Paper. Available online: https://www.bis.org/publ/mktc07.pdf (accessed on 1 October 2024).

Bloomfield, Robert, and Maureen O'Hara. 1999. Market transparency: Who wins and who loses? *The Review of Financial Studies* 12: 5–35. [CrossRef]

Bloomfield, Robert, Maureen O'hara, and Gideon Saar. 2005. The "make or take" decision in an electronic market: Evidence on the evolution of liquidity. *Journal of Financial Economics* 75: 165–99. [CrossRef]

Chelley-Steeley, Patricia, Brian Kluger, James Steeley, and Paul Adams. 2015. Trading Patterns and Market Integration in Overlapping Experimental Asset Markets. *Journal of Financial and Quantitative Analysis* 50: 1473–99. [CrossRef]

Dunne, Peter, Michael Moore, and Richard Portes. 2006. European government bond markets: Transparency, liquidity, efficiency. In *CEPR Working Paper*. London: Centre for Economic Policy Research.

Fischbacher, Urs. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10: 171–78. [CrossRef]

Flood, Mark D., Ronald Huisman, Kees G. Koedijk, and Ronald J. Mahieu. 1999. Quote disclosure and price discovery in multiple-dealer financial markets. *The Review of Financial Studies* 12: 37–59. [CrossRef]

Foucault, Thierry, Sophie Moinas, and Erik Theissen. 2007. Does anonymity matter in electronic limit order markets? *The Review of Financial Studies* 20: 1707–47. [CrossRef]

Gozluklu, Arie E. 2016. Pre-trade transparency and informed trading: Experimental evidence on undisclosed orders. *Journal of Financial Markets* 28: 91–115. [CrossRef]

Greiner, Ben. 2004. *An Online Recruitment System for Economic Experiments*. Gottingen: GWDG Bericht.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *The American Economic Review* 70: 393–408.

Majois, Christophe. 2008. Anonymity in a Limit-Order Market: An Experimental Analysis. Working Paper. Available online: https://www.semanticscholar.org/paper/Anonymity-in-a-Limit-Order-Market%3A-An-Experimental-Majois/407dbdfcd9c85dc60757457949b5a5b3f2411c55 (accessed on 1 October 2024).

Milgrom, Paul, and Nancy Stokey. 1982. Information, trade and common knowledge. *Journal of Economic Theory* 26: 17–27. [CrossRef]

O'Hara, Maureen, and Xing Alex Zhou. 2021. The electronic evolution of corporate bond dealers. *Journal of Financial Economics* 140: 368–90. [CrossRef]

Perotti, Pietro, and Barbara Rindi. 2006. Market for information and identity disclosure in an experimental open limit order book. *Economic Notes* 35: 97–119. [CrossRef]

Reiss, Peter C., and Ingrid M. Werner. 2005. Anonymity, adverse selection, and the sorting of interdealer trades. *The Review of Financial Studies* 18: 599–636. [CrossRef]

Rindi, Barbara. 2008. Informed traders as liquidity providers: Anonymity, liquidity and price formation. *Review of Finance* 12: 497–532. [CrossRef]

Ruiz-Buforn, Alba, Simone Alfarano, Eva Camacho-Cuena, and Andrea Morone. 2021. Single vs. multiple disclosures in an experimental asset market with information acquisition. *The European Journal of Finance* 28: 1513–39. [CrossRef]

Simaan, Yusif, Daniel G. Weaver, and David K. Whitcomb. 2003. Market maker quotation behavior and pretrade transparency. *The Journal of Finance* 58: 1247–67. [CrossRef]