Journal of Financial Markets 59 (2022) 100732

Contents lists available at ScienceDirect

Journal of Financial Markets

journal homepage: www.elsevier.com/locate/finmar

Transparency in fragmented markets: Experimental evidence

Terrence Hendershott [a,*], Marvin Wee [b], Yuanji Wen [c]

[a] Haas School of Business, University of California, Berkeley, CA, 94720-1900, USA
[b] College of Business and Economics, The Australian National University, Canberra, ACT, 2601, Australia
[c] UWA Business School, University of Western Australia, Address: 35 Stirling Highway, Crawley, 6009, WA, Australia

A R T I C L E I N F O

JEL classification:
C92
G14
G24
G28

Keywords:
Hidden orders
Iceberg orders
Dark pool trading
Limit order book market
Laboratory test

A B S T R A C T

We experimentally examine pre-trade transparency in fragmented limit order markets. Allowing traders to hide their orders encourages limit order usage. This improves measures of liquidity by increasing depth and narrowing spreads. However, because some of this depth is not displayed, market fragmentation may limit traders' ability to capitalize on the improved liquidity. This happens when traders execute against orders at worse prices than orders in another market, often referred to as "trade-throughs." In our laboratory setting, increased trade-throughs in a dark market impose costs of similar magnitude to the benefits of increases in depth leaving effective liquidity unchanged.

1. Introduction

Market centers compete not only on price but also on the speed of execution and pre-trade transparency. While competition among stock exchanges and other trading centers can be beneficial by encouraging markets to meet the needs of traders and reducing trading fees, the proliferation of venues creates challenges for regulators. For instance, the U.S. Securities and Exchange Commission (SEC) has expressed concern that the lack of transparency can potentially undermine the quality of displayed prices (White, 2014).[1] In this paper, we provide evidence on the effects of opacity in fragmented markets in an experimental setting.

Market structure has evolved dramatically. Improvements in trading technology enable trading on many venues including dark pools, where trading occurs without pre-trade transparency. In the United States, there are 15 registered stock exchanges and 33 alternative trading systems (SEC, 2020). Such trading systems include broker-dealers' crossing systems that match their clients' (and sometimes their own) orders continuously without disclosing orders before trading occurs. Stock exchanges also offer nondisplayed orders that allow traders to hide their entire order (hidden order) or display only a fraction of the order (iceberg or reserve order).[2] In the 2010s, dark trading in the U.S. increased from 35% to 55% of trading volume.[3]

* Corresponding author.
E-mail addresses: hender@haas.berkeley.edu (T. Hendershott), marvin.wee@anu.edu.au (M. Wee), yuanji.wen@uwa.edu.au (Y. Wen).

[1] https://www.sec.gov/news/speech/2014-spch060514mjw.

[2] When there is no minimum display requirement, an iceberg order can become a hidden order. Some stock exchanges impose a minimum display requirement while others, such as the Nasdaq, do not. Brokers also offer synthetic reserve orders to their clients. Additionally, the peak size of an iceberg order could be set as a fixed number (characterized in the current study) or as a random amount falling within a pre-set band.

[3] See https://www.sec.gov/marketstructure/datavis/ma_exchange_hiddenvolume.html on exchange hidden volume and SEC (2013) and SEC (2020) for dark pool volumes.

https://doi.org/10.1016/j.finmar.2022.100732
Received 8 February 2021; Received in revised form 24 March 2022; Accepted 24 March 2022
Available online 31 March 2022
1386-4181/

New opaque trading venues such as dark pools typically increase both opacity and market fragmentation. Hence, disentangling the impact of fragmentation and opacity empirically is often challenging. The microstructure literature modeling traders' order choice (i.e., market vs. limit orders) has focused on the tradeoff between transaction costs (i.e., price improvement) and execution probability (Peterson and Sirri, 2002). Studies on opacity extend the consideration of order placement choice to incorporate the option to limit pre-trade transparency (e.g., Boulatov and George, 2013; Zhu, 2014; Bloomfield et al., 2015; Buti et al., 2017; Bayona et al., 2021).

When studying the effects of opacity on informed trading strategies, previous studies have not incorporated the effects of fragmentation and the costs associated with trade-throughs. A trade-through is a purchase in one market at a price higher than the best (lowest) available offer in another market, or a sale at a price lower than the best (highest) available bid. In consolidated markets with fully hidden or partially hidden (iceberg) orders, there are no trade-throughs as price-transparency-time priority ensures that a hidden limit order will be executed after all displayed depth at the same price, but before disclosed orders at worse prices. Foucault and Menkveld (2008) note that this is no longer the case in fragmented markets and that orders can be traded-through. The cost of a trade-through is imposed on the trader who received the inferior price and the trader whose bid or offer is traded-through, and is indicative of economically inefficient trades (Battalio et al., 2004).[4]

Trade-throughs can occur because traders are unaware of better prices available due to low pre-trade transparency (Harris, 2015) or because traders try to economize on monitoring costs and the effort required to split their orders (Foucault and Menkveld, 2008). To prevent trade-throughs, some regulators impose rules such as those by the SEC (Hendershott and Jones, 2005). Nondisplayed orders increase the difficulty of knowing the best price, which should increase the incidence of trade-throughs. Hence, we expect that trade-throughs are more frequent in markets with pre-trade opacity than those without.

Trade-throughs are likely to moderate the positive effects of opacity in a centralized market due to informed traders providing more liquidity. As trade-throughs can lower execution probabilities and reduce the profits from providing liquidity, informed traders are less likely to supply liquidity in venues where trade-throughs are more prevalent. Thus, the positive effect of opacity on competition in liquidity provision is likely weakened with trade-throughs in fragmented markets. If the costs of trade-throughs are less than the benefits of opacity, we predict greater liquidity provision by informed traders to increase the true depth (the sum of lit and hidden depth) and to narrow the true spread (including hidden orders).[5] However, the substitution of displayed orders for hidden orders can decrease displayed depth and widen displayed spread.

When considering the effective transaction costs traders incur, it is important to include the costs associated with trade-throughs as the costs of failing to always receive the best price can offset improvements in market liquidity that can arise from opacity. A decrease in transaction cost only occurs if liquidity demanders actually access the better prices/spread. Given that the effects of opacity and fragmentation on transaction costs are contingent on several factors, there is no clear directional effect of opacity on the effective spreads based on the prices traders actually receive as opposed to the liquidity that is available, but may be traded through.

Our experimental approach isolates the effects of opacity on trader and market behavior while holding fragmentation constant. By using a market with two identical lit trading venues as the benchmark, we can compare the market quality of opaque market types (i.e., hidden and dark) against this benchmark while maintaining the same level of market fragmentation. As in the case with many financial markets, our trading venues operate as limit order books. Our opaque venues also operate as limit order books, which is the most prevalent trading mechanism in U.S. dark pools.[6] Using our experimental setting with multiple-venue markets, we investigate the effects of pre-trade opacity on trading strategies and various measures of market quality including spread, depth, trade-throughs, and realized transaction costs.

Consistent with the above discussion, in our experiments we find that there are more trade-throughs in the dark market. We also find informed traders provide more liquidity in a market where there is greater pre-trade opacity. Liquidity, as measured by true depth and spreads, improves with opacity but the displayed depth and spread worsens. These findings on trading behavior are similar to Bloomfield et al. (2015), henceforth BOS, who examined the effects of opacity on trading strategies and market outcomes by varying opacity in a single venue in an experimental setting.[7] As in Boulatov and George's (2013) model, BOS (2015) find traders substitute hidden orders for displayed orders when they introduce the ability to hide. However, BOS (2015) do not find significant changes to market outcomes, such as liquidity and informational efficiency.

While opacity increases informed traders' propensity to use limit orders, it is also important to consider the value of their private information. In their study of the make or take decision in an electronic limit order book, Bloomfield et al. (2005) show that informed traders are more likely to use market orders when their information is more valuable. This is due to the informed traders wanting to

[4] In corporate bond markets with broker-dealers, trade-throughs can occur when the broker-dealer imposes a markup to compensate themselves for arranging the trade (Harris and Mehta, 2020).

[5] The true spread is the difference between the best bid and ask regardless of whether these are displayed or hidden. The displayed and the true spreads are identical in the visible market, but they can differ in the markets that allow hidden orders or with a dark pool.

[6] There are two main classes of dark pool pricing mechanisms: (1) dark pools that cross orders at the midpoint of national best bid and best order (NBBO), and (2) dark pools that allow price flexibility. The latter operates as nondisplayed limit order books, in which the execution price is determined by the supply and demand. Table 1 in Menkveld et al. (2017) shows the market shares of different types of dark venues as a fraction of total trading volume in their sample of U.S. dark venues that report to the Nasdaq trade-reporting facility. They find the market share of dark pools that allow flexibility in their execution price (i.e., not midpoint) is three times the market share of dark pools that use midpoint crossing. Using weekly volume trade data reported to the Financial Industry Regulatory Authority (FINRA), Brolley (2020) shows that 66.1% of dark pools accept limit orders that allow the trader to set a limit price and for the order to be filled at the prespecified price (or better).

[7] The three markets they study are: (1) visible markets, (2) iceberg markets that allow both displayed and partially displayed orders, and (3) hidden markets where orders can be displayed, partially displayed, or completely hidden.

Table 1

Game specifications and experimental design. This table illustrates the experimental design. Panel A presents the specifications of the 16 games/securities with market type defined in the text. More details are provided in Table A in Appendix 1. Panel B lists the sequence of these 16 games in each experiment. We ran each experiment twice with two different cohorts, giving us 12 experimental sessions in total.

Panel A: Game Specifications

GameID	Market Type	Settings[a] (Extremity_Dispersion)	True Value
(1)	LIT	Low_Low	36
(2)	LIT	Low_High	58
(3)	LIT	High_Low	82
(4)	LIT	High_High	21
(5)	LIT_LIT	Low_Low	47
(6)	LIT_LIT	Low_High	36
(7)	LIT_LIT	High_Low	86
(8)	LIT_LIT	High_High	32
(9)	LIT_HID	Low_Low	51
(10)	LIT_HID	Low_High	62
(11)	LIT_HID	High_Low	18
(12)	LIT_HID	High_High	68
(13)	LIT_DRK	Low_Low	54
(14)	LIT_DRK	Low_High	43
(15)	LIT_DRK	High_Low	23
(16)	LIT_DRK	High_High	74

Panel B: Experimental Design

Experiment No.	1	2	3	4	5	6
Sequence of games	(1)	(1)	(1)	(1)	(1)	(1)
	(2)	(2)	(2)	(2)	(2)	(2)
	(3)	(3)	(3)	(3)	(3)	(3)
	(4)	(4)	(4)	(4)	(4)	(4)
	(5)	(5)	(9)	(9)	(13)	(13)
	(6)	(6)	(10)	(10)	(14)	(14)
	(7)	(7)	(11)	(11)	(15)	(15)
	(8)	(8)	(12)	(12)	(16)	(16)
	(9)	(13)	(5)	(13)	(5)	(9)
	(10)	(14)	(6)	(14)	(6)	(10)
	(11)	(15)	(7)	(15)	(7)	(11)
	(12)	(16)	(8)	(16)	(8)	(12)
	(13)	(9)	(13)	(5)	(9)	(5)
	(14)	(10)	(14)	(6)	(10)	(6)
	(15)	(11)	(15)	(7)	(11)	(7)
	(16)	(12)	(16)	(8)	(12)	(8)

[a] If the true value is less than $17 away from the expected value of [0,100], i.e., $50, the parameter Extremity is low otherwise high. If the noise of signal informed traders receive is 2 as opposed to 10, the parameter Dispersion is low, indicating a high-precision signal otherwise high.

maximize the profit from their private information. Consistent with this, we show informed traders' likelihood to use limit orders decreases with the value of their private information.

While quoted market liquidity improves, neither liquidity traders nor informed traders in our experiments benefit in terms of transaction costs and profits as they fail to sufficiently access the better prices. This is because volume is much lower in the dark venue than in the benchmark lit markets. In addition, traders often trade against displayed limit orders in one venue while better priced non-displayed orders exist on the other venue (i.e., the hidden order is traded-through).[8] Overall, consistent with the mixed findings of the impact of dark trading in Bayona et al. (2021), the benefit of increased hidden liquidity is offset by traders' failure to access it due to trade-throughs.[9]

[8] Examining dark pools with the midpoint pricing mechanism, Aquilina et al. (2016) find dark orders can trade at stale prices that do not match the primary market midpoint due to latency issues. While trade-throughs that occurred in our experimental setting could be due to latency issues, traders' inability to observe the better prices in the hidden or dark venues likely plays a more significant role in the trade-throughs that we observe. This is because we observe more trade-throughs on the fully opaque venue.

[9] The use of sophisticated order routing to check all dark pools could potentially eliminate trade-throughs, but Anand et al. (2021) find that doing so is associated with higher costs.

While a number of our findings are consistent with BOS (2015), an important difference is that our multi-venue setting allows the study of the effects of transparency, together with trade-throughs on trading strategies and market outcomes. As noted by Foucault and Menkveld (2008), traders are reluctant to provide liquidity if limit orders are not protected from trade-throughs. We find evidence consistent with this.

In addition to the dark market, we also examine the effects of introducing a market offering the option to submit partially nondisplayed orders (henceforth hidden market). Traders make limited use of nondisplayed shares as roughly 10% of depth in the hidden market is not displayed. Thus, in contrast with the single-market findings in BOS (2015), we find that the use of hidden orders is lower in fragmented markets. In addition, our evidence on the effects of hidden and dark trading suggests that the form of pre-trade opacity in the presence of market fragmentation matters for overall market liquidity.

Most previous studies on market transparency examine the above two forms of opaque trading separately and without market fragmentation. Research on dark pools empirically shows that dark limit order book trading reduces spreads and increases informational efficiency (Foley and Putnins, 2016). However, price efficiency only improves with a low level of dark pool trading, not with a high level of dark pool trading (Comerton-Forde and Putnins, 2015). The relatively few studies examining nondisplayed orders finds price discovery and market quality improve when traders are able to completely or partially hide their orders (Kovaleva and Iori, 2015; Gozluklu, 2016). Linking these two branches of the literature, we compare the effect of the two forms of pre-trade opacity in fragmented markets on traders' order placement strategies and market quality. Our study provides insights that contribute to the ongoing debate over pre-trade transparency by documenting the positive and negative effects of opacity in fragmented markets and showing different market outcomes for the two types of opacity. Allowing for market fragmentation enables us to examine the interaction between transparency and trade-throughs to document trading behavior that has not been examined in previous studies.

2. The experiment

2.1. *Trading procedure, trader type, and information structure*

We conduct a series of in-person experimental asset trading sessions to investigate the effects of transparency on market quality. In each of the twelve sessions, a cohort of eight participants traded 16 securities sequentially for a period of 3 min each. These 16 securities differ in their true value and trading environment (i.e., market type and signal dispersion). The list of securities is provided in Panel A of Table 1 and explained below. In our analyses, we treat each cohort-security as one observation, unless specified otherwise. We use the term "venue" to indicate where a security is traded (i.e., lit, hidden or dark venue) and the term "market" is the collective term for the trading venues (e.g., LIT_DRK refers to the market with two venues, one lit and the other dark).

There are four informed traders and four liquidity traders in each trading period. Both informed traders and liquidity traders have the same endowment of cash and shares for each security and the endowment is not rolled over across the 16 securities. Each participant in the experimental session trades half of the securities as an informed trader and the other half as a liquidity trader. Following BOS (2015), in each trading period, two informed traders observe the sum of the true liquidating dividend plus a predetermined random number, while the other two observe the sum of the true liquidating dividend minus the same predetermined random number for that security. This information structure ensures that each informed trader has noisy information about securities, but the informed traders have precise information in the aggregate. The four liquidity traders are split evenly into two groups: buyers and sellers. One of the buyers (sellers) is tasked to buy (sell) 30 shares and the other to buy (sell) 40 shares and penalties are imposed on those who fail to meet their targets. The net demand of the four liquidity traders is zero.[10] Unlike informed traders, liquidity traders do not receive any private information about the true liquidating dividend and are told that the liquidating value is drawn from a bell-shaped distribution between 0 and 100. Introducing liquidity traders avoids the possibility of a no-trade equilibrium under extreme risk aversion condition. This is commensurate with most canonical theoretical models, such as Kyle (1985).

Like BOS (2015), we also manipulate both the dispersion of information provided to informed traders and the extremity of the security's value. Both parameters are used to influence the information asymmetry environment. In particular, the dispersion of information characterizes the precision of the signal received by informed traders. Its value is either $2 (low dispersion setting) or $10 (high dispersion setting). For example, if the dispersion is $10, we draw a random number, say y, between the signal range $[-10, 10]$. Two informed traders are told that the range of true liquidating dividend is within (True Value – y – 10, True Value – y + 10) while the other two are told (True Value + y – 10, True Value + y +10). Note that the dispersion directly determines the range of the informed traders' information (i.e., 2 × dispersion parameter) and it also affects the expected value of the ranges each group of informed traders receive, i.e., (True Value – y) and (True Value + y) respectively. All traders including the liquidity traders are told which dispersion setting they face in each trading period. The smaller the value of the dispersion for a security, ceteris paribus, the higher the information asymmetry environment is between informed and liquidity traders.

In addition, extremity affects the value of the information to the informed traders and the degree of information asymmetry among the traders. Imagine that all participants are given the unconditional distribution and that the true liquidating divided is significantly

[10] This is the same as Gozluklu (2016) but different from BOS (2015), who has either a positive or negative net demand in their experimental sessions (Table 1, Panel B, p. 2239). There are two reasons for our choice. First, the non-zero net demand condition can result in one-sided price pressure and add noise to the market outcome. Second, as explained in Subsection 2.1, we run a setting of a single lit market in addition to three market types that we focus on. Given that the full factorial combination of all controlled parameters, market type (4) × extremity (2) × dispersion (2), is 16, we could not take the time to consider another factor (net demand) in the 2-h experimental session.

larger or smaller than the prior unconditional expected value of $50. If an informed trader correctly infers the true liquidating dividend from their private information, they have a good chance of making relatively large profits. Hence, the information asymmetry is higher in a high extremity setting than in a low extremity setting. In our analysis, we define the extremity setting as high when the liquidating dividend deviates by at least $17 from the prior unconditional value of $50, the extremity setting is low when the liquidating dividend deviates no more than $16.

Our experimental design manipulates the following three factors: market type (LIT, LIT_LIT, LIT_HID, LIT_DRK, explained below), extremity (Low, High), and dispersion (Low, High). The corresponding full factorial $4 \times 2 \times 2$ within-subjects design is presented in Table 1. Each security features one of these 16 combinations. To account for potential sequence effect, we change the order of securities and define six experiments shown in Panel B, Table 1. First, since our primary focus is market type, we change the order of the blocks defined by market type (a block means four securities associated with the factorial 2 extremity settings $\times$ 2 dispersion settings combinations) while maintaining the sequence within each block. Second, we ensure the "LIT" block is played first as part of the training. In the analysis below, we do not include a discussion on the LIT market for brevity.

2.2. Market types

Our experimental asset markets are anonymous limit order book markets (continuous-time double auction) where the execution mechanism is similar to that widely used in many global stock markets. The price-visibility-time execution priority rule is enforced in our markets. With all markets, the activities are anonymous and participants cannot identify, other than themselves, who is associated with a certain order. Past trades (i.e., volume and price traded) are reported to all participants but the identity of those involved are not reported. The markets differ by the availability of different order types (i.e., limit orders, market orders, and nondisplayed orders). We examine the following market types: (1) two identical lit trading venues (LIT_LIT) simultaneously open for trading, (2) two lit venues simultaneously open for trading with one venue allowing traders to submit partially and fully nondisplayed orders (LIT_HID), and (3) two trading venues where one is lit and the other is dark (LIT_DRK). All markets have a tick size of $1.

In the LIT_LIT market, the two trading venues have identical features where traders are free to submit market orders or limit orders, and can choose to amend or cancel unexecuted limit orders that they have previously submitted.[11] To avoid any potential labelling effects that might arise from unobserved preferences of the participants, we use abstract symbols instead of numbers, alphabets or words to label the venues (e.g., Costa-Gomes et al., 2001). For instance, ## and @@ are used to tag the two venues in the LIT_LIT market.

In the LIT_HID market, traders have the option of partially or fully hiding their order in one of these two venues (denoted as the HID venue). In the HID venue, traders may choose to submit limit orders, market orders or iceberg orders. When selecting to submit an iceberg order, the trader has to specify the units to display. If the trader chooses to display zero units, the entire order is hidden. Similar to how iceberg orders work in practice, the nondisplayed shares have lower execution priority over displayed shares at a given price.[12]

In the last market type, the LIT_DRK market, the limit order book of one of the two venues is not observable (denoted as the DRK venue). Similar to the LIT venue, traders in the DRK venue can submit limit orders and market orders, and the execution mechanism is the same across the two venues. These functional features characterize a dark limit order book market.[13]

These three market types present a rich setting to study the issues associated with pre-trade opacity. Using the LIT_LIT market as our control group, comparisons between the control group and the LIT_DRK markets help us to understand how exogenous opacity contributes to market quality and informational efficiency in the context of a dual-venue framework. Similarly, by examining the market outcome of the LIT_LIT and the LIT_HID markets, we can study the effect of endogenous opacity on market quality and extend the findings in BOS (2015) to a dual-venue framework. Last, the comparison between the LIT_HID and the LIT_DRK markets allows us to directly compare whether these two forms of opacity differ. Sample screens of these markets are provided in Appendix 1 to illustrate the trading platform used in the experiments.

2.3. Subjects, instructions, and incentives

We ran all experimental sessions in September 2018. University students were invited to a training session before participating in the formal session, with a maximum of 18 days between the two sessions for any individual. The participants in the study are students from both undergraduate (65%) and graduate programs (35%) with majors in accounting (27%), economics (4%), finance (60%), and others (9%). Gender is balanced with 47 female students and 49 male students. The activities in each training session were organized as follows. Upon their arrival in the room, each participant received a set of detailed experimental instructions and a randomly assigned log-in token. When all registered participants were present, the experiment administrators went through the instructions and answered

[11] On the screen layout, the information of one venue is placed above the other in the Market View and Market Depth windows. To avoid any display effects, we use two screen layouts for each security. The two screen layouts differ by which venue is displayed on the top and bottom halves of the windows. The two screen layouts are randomly assigned to the eight participants in the game.

[12] This is a key feature of iceberg orders or reserve orders. Examples on how iceberg orders are executed can be found in the manual published by the London Stock Exchange. See https://www.londonstockexchange.com/products-and-services/technical-library/technical-guidance-notes/technicalguidancenotesarchive/setsmm-and-iceberg.pdf.

[13] We adopt features that are mostly prevailing in U.S. dark pools, as opposed to mid-execution trading mechanism, to isolate the effects of opacity from the trading mechanism.

any questions. The participants were then instructed to complete the 16 trading games.

At the start of the formal sessions, participants were asked to review the experimental instructions and complete an online quiz. which provided immediate feedback.[14] This ensured the participants understood the incentives and the games in the experiment. Once all participants had completed their quiz, the administrators instructed the participants to complete the trading games. Upon completing the games, participants were asked to complete a short demographic survey before they were provided with their cash payment. We paid participants for participating in the formal sessions and all analyses are based on the data collected from these sessions. Each participant was paid \$30 ± \$5 for every 1000 laboratory dollars gained or lost through trading and penalties, to a minimum of \$10 dollars. We told the participants the formula used to compute their winnings. When participating in the games as a liquidity trader, a penalty was imposed if they did not meet their liquidity obligation. Specifically, each trader received a 100 laboratory dollar penalty for each share that they fell short of their target.

3. Empirical strategies and results

We employ repeated-measures ANOVA tests (also called within-subjects ANOVA) to analyze the data from the 12 experimental sessions (with 12 cohorts of subjects). We do so because our experiment uses a within-subjects design to investigate the market outcome of different market conditions and hence observations of market conditions under investigation are from related, not independent groups. In general, for analyses of individual-level variables (such as number of limit orders), we compute these variables for the average trader (of a certain informational role, if needed) in a cohort. For analyses of market-level or venue-level variables (such as total depth), we first define each 30 seconds as one interval and use time as weights within the 30-s interval to obtain an average observation for the interval. These interval-level observations are then used in the two-way repeated measures ANOVA tests where market type and interval are treated as two dimensions for statistical testing. In the analysis of transaction cost, we compute effective spread as well as its components in the following way: we compute these variables trade-by-trade, then take the simple average to obtain an observation for each cohort per security.[15] Then one-way ANOVA tests are adopted to test the differences across market type.

Table 2 presents summary statistics on the orders and cancellations by market type. For all market types, there are generally more limit orders placed, followed by market orders (market and marketable limit orders), and amendments/cancellations. As shown in Panel A, the median limit order size is 10 shares across all three market types. In contrast, the mean order size differs across the three market types and there is heterogeneity in the size of limit orders placed in each market type. For instance, the 5th percentile order size in the LIT_LIT is five shares and the 95th percentile is 40 shares. We observe similar distributions with the market orders in Panel B and amendments and cancellations in Panel C.

Table 2 shows larger limit orders and more cancellations are used in the LIT_DRK market. This is consistent with traders responding to lower execution probabilities in the dark venue. This together with variations in order sizes suggest that analytical models should incorporate these features when investigating the effects of transparency and fragmentation. Because doing so has proven challenging due to the tractability of the models, an experimental setting can provide insights that models have struggled to provide.

Table 2
Orders and cancellations by market type. This table presents the distribution of the sizes as measured by number of shares of limit orders (Panel A), market orders (Panel B), and amendments and cancellations (Panel C) by market type.

	n	Mean	Std Dev	5th percentile	25th percentile	Median	75th percentile	95th percentile
Panel A: Limit orders								
LIT_LIT	2736	16.67	13.81	5	10	10	20	40
LIT_HID	3192	14.62	13.82	1	5	10	20	40
LIT_DRK	3056	18.50	15.99	5	10	10	20	50
Panel B: Market orders								
LIT_LIT	875	14.27	15.51	1	5	10	20	40
LIT_HID	1031	12.64	11.11	1	5	10	15	40
LIT_DRK	898	12.97	10.91	1	10	10	20	30
Panel C: Amendments/cancellations								
LIT_LIT	631	20.79	16.66	10	10	20	30	50
LIT_HID	663	20.33	14.04	5	10	15	30	50
LIT_DRK	839	23.23	16.89	5	10	20	30	50

[14] To ensure the participants are familiar with the setup, they were allowed to access the quiz as many times as they wanted. They could also ask the experiment administrators any questions during this stage.

[15] The results using share-weighted measures are qualitatively similar and are omitted for brevity.

Table 3

Trade-through statistics. This table presents the trade-through measures for informed and liquidity traders in the LIT_LIT, LIT_HID, and LIT_DRK markets. *TT* is defined as the number of trades that occurred at prices worse than the best "true" bid and ask prices ("traded-through" trades), scaled by the total number of trades initiated by the corresponding informational role in a given market. *TTshr* is defined as the number of shares traded-through scaled by the total number of shares in trades initiated by the corresponding informational role in a given market. The measures are averages based on 48 observations (i.e., 12 cohorts of 4 securities for each market type). The *p*-values for "Informed against Liquidity" are based on *t*-tests between the informed and liquidity trader groups, across the three market types. The repeated measures ANOVA tests are based on 192 market-trader type level observations, where trade-through for each trader type in each market is examined separately.

		LIT_LIT (1)	LIT_HID (2)	LIT_DRK (3)	*t*-test Informed against Liquidity (*p*-value)	ANOVA test (*p*-value)		
						(1) = (2)	(1) = (3)	(2) = (3)
TT	Informed	0.11	0.13	0.19	0.11	0.08	0.00	0.02
	Liquidity	0.11	0.17	0.26				
TTshr	Informed	0.11	0.13	0.17	0.02	0.30	0.00	0.06
	Liquidity	0.14	0.18	0.24				

3.1. Effects of pre-trade opacity on trade-throughs

As discussed in the Introduction, trade-throughs are trades that occur at a non-optimal price and may arise in fragmented markets for several reasons.[16] The ability to consider the occurrence and effects of trade-throughs is a pertinent difference between our study and BOS (2015). We predict trade-through are more frequent when markets are pre-trade opaque. To demonstrate that traders are missing out on opportunities to trade at the best quotes available, in Table 3 we examine the trade-through activities by market type. We also examine trade-throughs by informational roles because informed traders have more information about the true value than liquidity traders and may be less likely to experience trade-throughs.[17] We consider two measures: *TT* is defined as the number of trades that occurred at non-optimal prices scaled by the total number of trades initiated by the corresponding informational role in a given market, while *TTshr* is defined as the number of shares of those trade-through trades scaled by the total number of shares initiated by the corresponding informational role in a given market.

Table 3 shows, when comparing across market types, there are significantly more trade-throughs in the LIT_DRK than the LIT_LIT market (*TT* *p*-value = 0.00, *TTshr* *p*-value = 0.00). Similarly, there are more trade-throughs in the LIT_DRK than in the LIT_HID market (*TT* *p*-value = 0.02, *TTshr* *p*-value = 0.06). However, the difference in the trade-through activities between the LIT_HID and LIT_LIT markets are only statistically different when examining the number of trade-throughs (*TT* *p*-value = 0.08). When we partition the trades by informational roles, we find trade-through volume measured as a proportion is significantly higher for liquidity traders (14%–24%) than informed traders (11%–17%), although less so when we measure trade-through activities by the number of trades (*TTshr* *p*-value = 0.02, *TTshr* *p*-value = 0.11). Overall, these findings help corroborate our predictions that trade-throughs are more frequent when markets are more pre-trade opaque.

3.2. Effects of pre-trade opacity on traders' order submission strategies

Of key interest to us is the extent to which traders' strategies are affected by their ability to hide orders in a multi-venue market. As shown in Fig. 1, the total number of limit order shares submitted by both informed and liquidity traders are higher in the LIT_DRK market than in the LIT_HID and LIT_LIT markets. The volume of limit orders submitted in the LIT_DRK market is also significantly higher than the other two market types (untabulated, three-market *p*-value = 0.00, two-market *p*-value = 0.01).

First, we examine whether there is an increase in traders' usage of nondisplayed orders with pre-trade opacity. On the LIT_HID market, 12% of limit orders submitted by informed traders are not displayed. While liquidity traders exhibit similar behavior, their usage of nondisplayed limit orders is lower at 6%. Compared to the LIT_HID market, traders in the LIT_DRK market have an even greater propensity to substitute from lit orders to nondisplayed orders. Of limit orders submitted by informed and liquidity traders, 45% and 43%, respectively, are dark.[18] The untabulated ANOVA test shows the volume of nondisplayed orders submitted in the LIT_DRK market is significantly higher than the volume of nondisplayed limit orders submitted in the LIT_HID market (*p*-value = 0.00). Overall, these observations suggest that traders substitute nondisplayed orders for displayed orders when the market allows pre-trade opacity. It is interesting to note the relatively low proportion of nondisplayed orders used in the LIT_HID market when compared to nondisplayed orders used in the LIT_DRK market.

[16] It could be because the quotes of some but not all trading venues are consolidated and, even if all quotes are consolidated, latency in data feed from these trading venues may also cause trade-throughs to occur.

[17] These trades can occur in our experiment because: (1) traders are unable to observe the "true" bid and ask prices available in the LIT_HID and LIT_DRK markets, (2) a limit order that improves the best bid or ask prices is placed in a co-existing venue at the same time that the trader executes her order in the other venue, or (3) traders fail to check the co-existing venues for the best venue to place their orders. A smart order router would eliminate the third scenario but not the first two.

[18] This observation is made at the market level. If we compare the usage of limit orders and market orders by informed and liquidity traders at the venue level, we find that traders favor the lit venue over the dark venue in the LIT_DRK market (the relevant *p*-values are all not larger than 0.01).

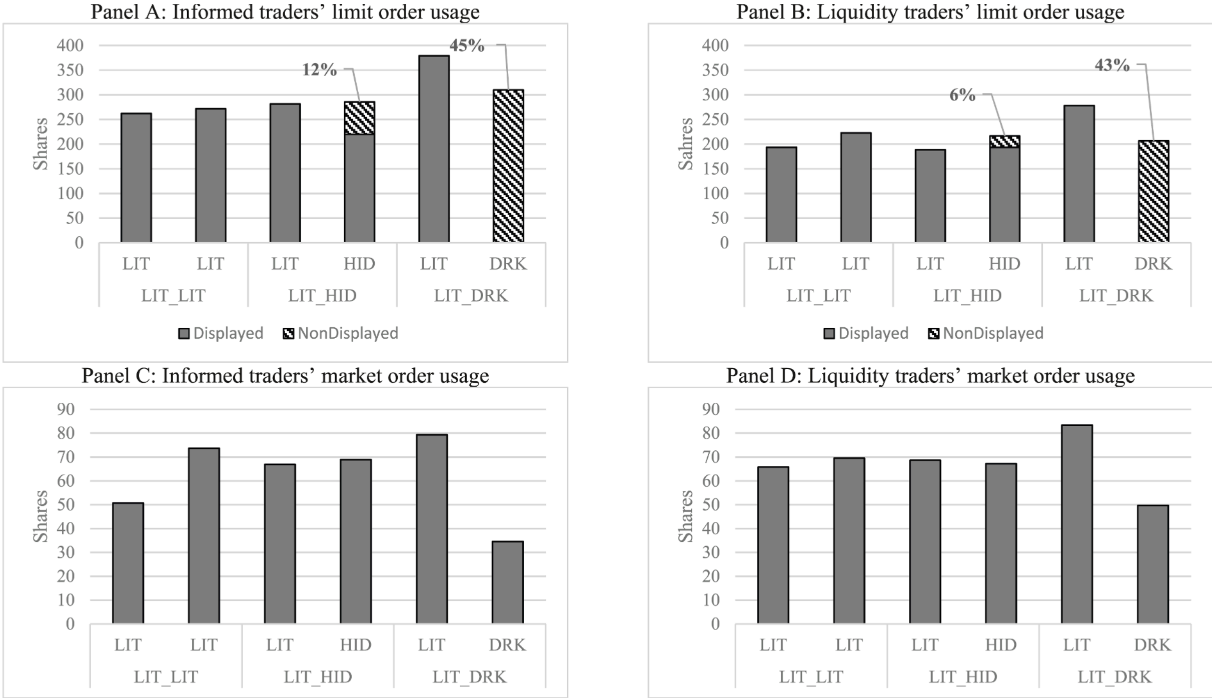

Fig. 1. Number of shares in limit and market orders submitted by informed/liquidity trader. This figure presents summary statistics on the submission of limit orders and market orders. In each panel, we consider the number of displayed and nondisplayed shares submitted by either informed or liquidity traders in the three markets: LIT_LIT, LIT_HID, and LIT_DRK. The measures plotted are averaged across the cohorts by venue for each market type. In Panels A and B, we plot the average number of limit order shares placed by informed traders and liquidity traders, respectively. In Panels C and Panel D, we plot the average number of market order shares placed by informed and liquidity traders respectively.

Using the measures studied in BOS (2015), we examine the make-take decisions by market type (Table 4), and by venue and trader type (Table 5). The three measures are submission rate (*SR*), fill rate (*FR*), and taking rate (*TR*). *SR* is defined as the number of shares in limit orders divided by the total number of limit and market order shares submitted and indicates a trader's propensity to provide liquidity. *FR*, defined as the number of executed shares in limit orders divided by the total number of shares submitted in limit orders, measures the extent that a limit order is filled. *TR*, defined as the number of shares a trader executed using market orders divided by the total number of shares she traded, measures how likely a trader trades using market orders.[19] We also decompose each of these three measures into displayed and nondisplayed components. For instance, *SR* is decomposed into *DSR* (defined as the number of displayed shares in limit orders over the total number of shares submitted) and *NDSR* (defined as the number of nondisplayed shares in limit orders over the total number of shares submitted).[20]

Table 4 shows that, among the three market types, traders have a higher propensity to provide liquidity in the LIT_DRK market. The submission rate (*SR*) is 0.82 in the LIT_DRK market, as compared to 0.78 in both the LIT_LIT and LIT_HID markets. The differences in liquidity provision between LIT_DRK and LIT_LIT, and between LIT_DRK and LIT_HID are statistically significant (p-values = 0.03 and 0.01, respectively). When comparing the effects of exogenous and endogenous opacity, we find only the former has a significant impact on liquidity provision.

The average fill rate (*FR*) for all traders in the LIT_DRK market is lower than in the LIT_HID (p-value = 0.03) but not different from the LIT_LIT. For the LIT_HID market, the *NDFR* is 1% because only a small proportion of limit orders are not displayed (see Fig. 1). This is lower than what BOS (2015) document for a single-venue hidden market (4%–8%), indicating the reduced benefits of using non-displayed orders in fragmented markets.[21] To allow comparison between the LIT_HID and LIT_DRK markets, we also compute the fill rates of fully hidden limit orders by dividing the total number of shares in fully hidden orders traded by the number of shares in fully

[19] As the *TR* examines all trades rather than orders, it is a subset of *SR*.

[20] Note that $SR = DSR + NDSR$.

[21] One important difference between our experimental design and that of BOS (2015) is that our experimental number of combinations precludes us from having an aggregate demand treatment as well. Hence, aggregate demand is zero in our setup (as in Gozluklu (2016)) versus −20 or +20 in BOS (2015). Aggregate demand being zero implies that order imbalances are more likely to arise from the informed traders than the liquidity traders. This implies that informed traders have a greater incentive to hide their orders to limit information leakage in our setting than in BOS (2015). Thus, the lower use of hidden orders by informed traders in our experiments is likely due to fragmentation and not differences in aggregate demand.

Table 4

Make-or-take decisions and market making by market type. This table presents the submission rate, fill rate, taking rate, and market making at the market level. Submission rate (*SR*) is defined as the number of shares in limit orders divided by the total number of shares the traders submit in both limit and market orders. As the number of shares in limit orders comprises displayed shares and non-displayed shares, *SR* consists of displayed submission rate (*DSR*) and non-displayed submission rate (*NDSR*), i.e., $SR = DSR + NDSR$. *DSR* (*NDSR*) is the number of displayed (nondisplayed) shares in limit orders divided by the number of shares in both limit and market orders. Fill rate (*FR*) is the number of executed shares in limit orders divided by the total number of shares submitted in limit orders. *FR* comprises displayed fill rate (*DFR*) and non-displayed fill rate (*NDFR*), i.e., $FR = DFR + NDFR$. Taking rate (*TR*) is defined as the number of shares traders executed using market orders divided by the total number of shares traded. *TR* comprises displayed taking rate (*DTR*) and non-displayed taking rate (*NDTR*). i.e., $TR = DTR + NDTR$. *DTR* (*NDTR*) is defined as the number of displayed (nondisplayed) shares traders executed using market orders divided by the total number of shares traded. *MMF* is the fraction of traders who have placed both buy and sell limit orders in the same game. *MMT* is the fraction of the duration of the game where a trader has both buy and sell limit orders in the order book. The measures presented are averaged across 48 observations (i.e., 12 cohorts of 4 securities for each market type). The repeated measures ANOVA tests are based on 96 market level observations.

	LIT_LIT (1)	LIT_HID (2)	LIT_DRK (3)	ANOVA test (*p*–value)		
				(1) = (2)	(1) = (3)	(2) = (3)
SR	0.78	0.78	0.82	0.81	0.02	0.01
DSR	0.78	0.71	0.46	0.00	0.00	0.00
NDSR	0.00	0.07	0.36	N/A	N/A	0.00
FR	0.21	0.22	0.19	0.32	0.26	0.03
DFR	0.21	0.21	0.13	0.67	0.00	0.00
NDFR	0.00	0.01	0.06	N/A	N/A	0.00
TR	0.50	0.50	0.50	N/A	N/A	N/A
DTR	0.50	0.49	0.34	N/A	N/A	0.00
NDTR	0.00	0.01	0.16	N/A	N/A	0.00
MMF	0.47	0.41	0.51	0.12	0.14	0.00
MMT	0.23	0.21	0.27	0.34	0.08	0.01

Table 5

Make-or-take decisions at the venue level within market type by trader type. This table presents the submission rate, fill rate, and taking rate by different informational roles at the venue. Submission rate (*SR*) is defined as the number of shares in limit orders divided by the total number of shares the traders submit in both limit and market orders at the venue. Fill rate (*FR*) is the number of executed shares in limit orders divided by the total number of shares submitted in limit orders at the venue. Taking rate (*TR*) is defined as the number of shares traders executed using market orders divided by the total number of shares traded at the venue. The measures presented are averaged across 48 observations (i.e., 12 cohorts of 4 securities for each market type). The repeated measures ANOVA tests are based on 192 venue-trader type level observations, where the make-take decision of each trader type in each venue is examined separately.

		LIT_LIT		LIT_HID		LIT_DRK		ANOVA test (*p*–value)		
		LIT	LIT	LIT	HID	LIT	DRK	(1)=(2)	(3)=(4)	(5)=(6)
		(1)	(2)	(3)	(4)	(5)	(6)			
SR	Informed	0.84	0.81	0.82	0.81	0.82	0.90	0.50	0.64	0.00
	Liquidity	0.74	0.75	0.73	0.75	0.77	0.83			
FR	Informed	0.19	0.22	0.21	0.20	0.20	0.15	0.96	0.79	0.00
	Liquidity	0.25	0.22	0.26	0.26	0.26	0.16			
TR	Informed	0.45	0.41	0.40	0.46	0.42	0.38	0.75	0.82	0.09
	Liquidity	0.52	0.55	0.54	0.50	0.47	0.42			

hidden limit orders submitted. The fill rate for fully hidden limit orders placed in the LIT_HID market is not different from nondisplayed orders placed in the LIT_DRK market (untabulated), and lower than for displayed orders. By definition, the *TR* for the market is 0.5. We note similarities in the displayed taking rate (*DTR*) for the LIT_LIT market and the LIT_HID market but significant differences between the LIT_HID and LIT_DRK markets (0.49 vs. 0.34, p-value $= 0.00$).

To further examine the liquidity provision by the traders, we investigate the role that they play as market makers by measuring the fraction of traders who have placed both buy and sell limit orders in the same game (*MMF*) and the fraction of the duration of the game

Table 6
Effects of information asymmetry on make-or-take decisions in the dark market. This table presents the submission rate and taking rate by different informational roles at the venue level (e.g., columns (1) and (2)) by the level of information asymmetry. Submission rate (*SR*) is defined as the number of shares in limit orders divided by the total number of shares the traders submit in both limit and market orders. Fill rate (*FR*) is the number of executed shares in limit orders divided by the total number of shares submitted in limit orders. Taking rate (*TR*) is defined as the number of shares traders executed using market orders divided by the total number of shares traded. Information asymmetry is varied by the level of extremity of the security value. The measures presented are averaged across 24 observations.

		Low information asymmetry		High information asymmetry		*t*-test (*p*-value)	
		LIT	DRK	LIT	DRK	(1) = (3)	(2) = (4)
		(1)	(2)	(3)	(4)		
SR	Informed	0.82	0.93	0.82	0.86	0.97	0.02
	Liquidity	0.77	0.81	0.76	0.84	0.80	0.55
FR	Informed	0.18	0.13	0.22	0.17	0.30	0.41
	Liquidity	0.28	0.14	0.25	0.17	0.42	0.51
TR	Informed	0.47	0.29	0.37	0.47	0.16	0.06
	Liquidity	0.44	0.52	0.51	0.31	0.26	0.02

where a trader has both buy and sell limit orders in the order book (*MMT*). The market making measures show a greater fraction of traders act as market makers (*MMF*) in the LIT_DRK market than in the LIT_HID market (two market p-value = 0.00) and for a greater fraction of the duration of the game (*MMT*) in the LIT_DRK market than the LIT_HID market (two market p-value = 0.01) and the LIT_LIT market (p-value = 0.08).[22]

Our results contrast somewhat with BOS's (2015) findings that informed traders provide more liquidity when they can use non-displayed orders in the hidden market. The difference in the results may be due to the trade-throughs in fragmented markets, as shown in Subsection 3.1. In a single market where nondisplayed orders are permitted, the benefit of hiding shares is offset by the loss of execution priority to the displayed orders at the same price. In a multi-venue market, the benefit is further offset because hidden orders are more likely to be traded through. The total costs associated with using hidden orders may be sufficiently severe in a dual-venue market such that nondisplayed orders are no longer attractive for informed traders when realizing informational profits.

In examining the *SR* at the venue level and by different trader type, we note from Table 5 that the use of limit orders to market orders (*SR*) in the dark venue are much higher than that in the co-existing LIT venue: 8% and 6% higher for informed and liquidity traders, respectively.[23] This supports our prediction that informed traders provide more liquidity in a venue where there is pre-trade opacity and suggests the costs of trade-throughs do not overwhelm the benefits of opacity. However, consistent with the prediction in Zhu (2014), the fill rate in the dark venue (average across trader types *FR* = 16%) is almost 8% lower than in the coexisting lit venue. The take rate in the dark venue is 4% and 5% lower than in the lit venue for informed and liquidity traders, respectively.

To further our understanding of how opacity affects informed traders' liquidity provision in the dark market, we examine the effect that information asymmetry has on order submission strategies. In the experiment, we manipulate the extremity of the realized value to vary the value of the information to the informed traders. Table 6 shows the submission rate (*SR*) for informed traders is significantly lower in the dark venue with high information asymmetry than with low information asymmetry (0.86 vs. 0.93, p-value = 0.02). The take rate (*TR*) for the informed trader is also significantly higher in the dark venue when information asymmetry is high (0.47 vs. 0.29, p-value = 0.06). This is consistent with the expectation (Bloomfield et al., 2005) that informed trader's demand for liquidity increases with the value of their information.[24] Our findings suggest a trade-off between profiting from liquidity provision and maximizing profit from the private information.

3.3. Effects of pre-trade opacity on spreads and depth

In this subsection, we examine the effects of opacity on depth and spreads. In our experiments, the limit order book is empty at the start of each 3-min period and depth builds over time. Fig. 2 depicts the evolution of depth and spreads. Panels A and B show the dynamics of the total depth (*TrueDepth*) and displayed depth (*DispDepth*) over the six trading intervals. There is a general increase for all markets, but the extent of the increases differs across market types. Panels C and D show that spread (*TrueSpread*) and displayed

[22] The ANOVA three-market test examines the differences across the three markets: LIT_LIT, LIT_HID, and LIT_DRK. The two-market test, unless specified, examines the difference between the LIT_HID and LIT_DRK markets.
[23] It is not surprising to find informed traders are more likely to provide liquidity than liquidity traders (e.g., 0.84 > 0.78 in the LIT_DRK market). Although earlier studies assume informed traders use market orders exclusively, recent empirical studies find limit orders are used by both informed and uninformed traders [observed and discussed also in BOS (2015)].
[24] In untabulated results, we also see a reduction in the duration that informed traders quote on both sides of the market in the LIT_DRK market when information asymmetry is higher.

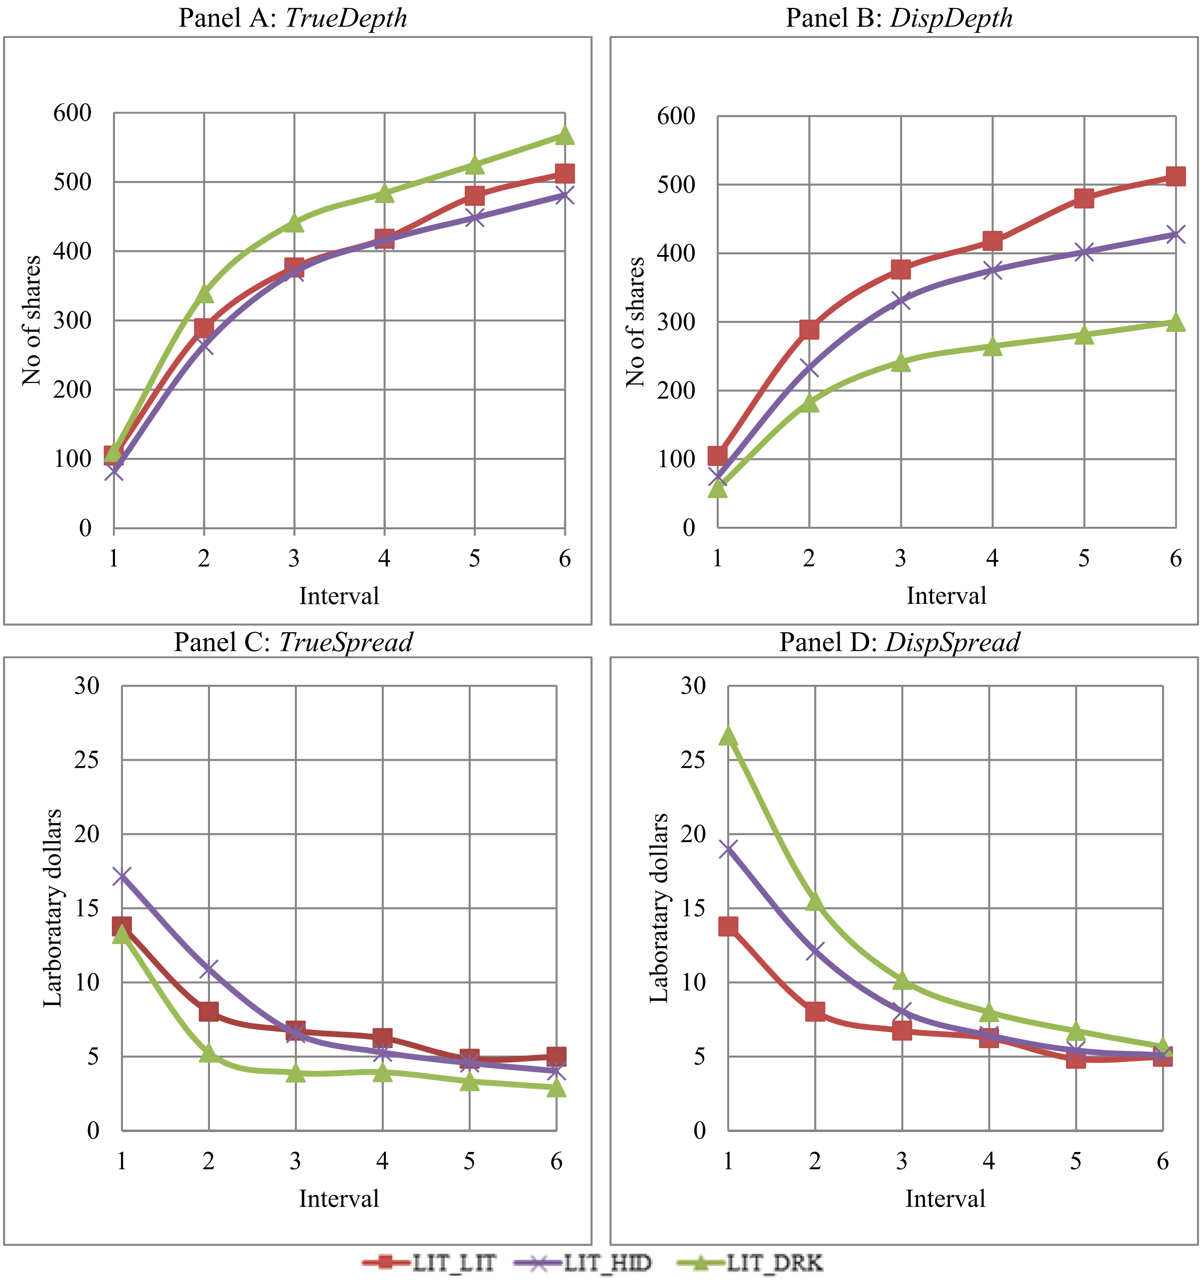

Fig. 2. Evolution of spread and cross-sectional comparison of depth and spread measures. This figure presents both the dynamics and the cross-market comparisons of liquidity during the 3-min trading game. The three minutes are partitioned into six 30-s intervals. Panels A and B present total true depth (*TrueDepth*) and displayed depth (*DispDepth*) over time, respectively. Panels C and D present total true spread (*TrueSpread*) and displayed spread (*DispSpread*) over time, respectively. *TrueDepth* is the number of both displayed and non-displayed shares at all price levels. *DispDepth* is the number of displayed shares at all price levels, while *TrueSpread* is the difference between the best "true" ask and best "true" bid, where "true" best bid and best ask prices is determined by examining both displayed and nondisplayed shares. *DispSpread* considers only displayed best bid and best ask prices.

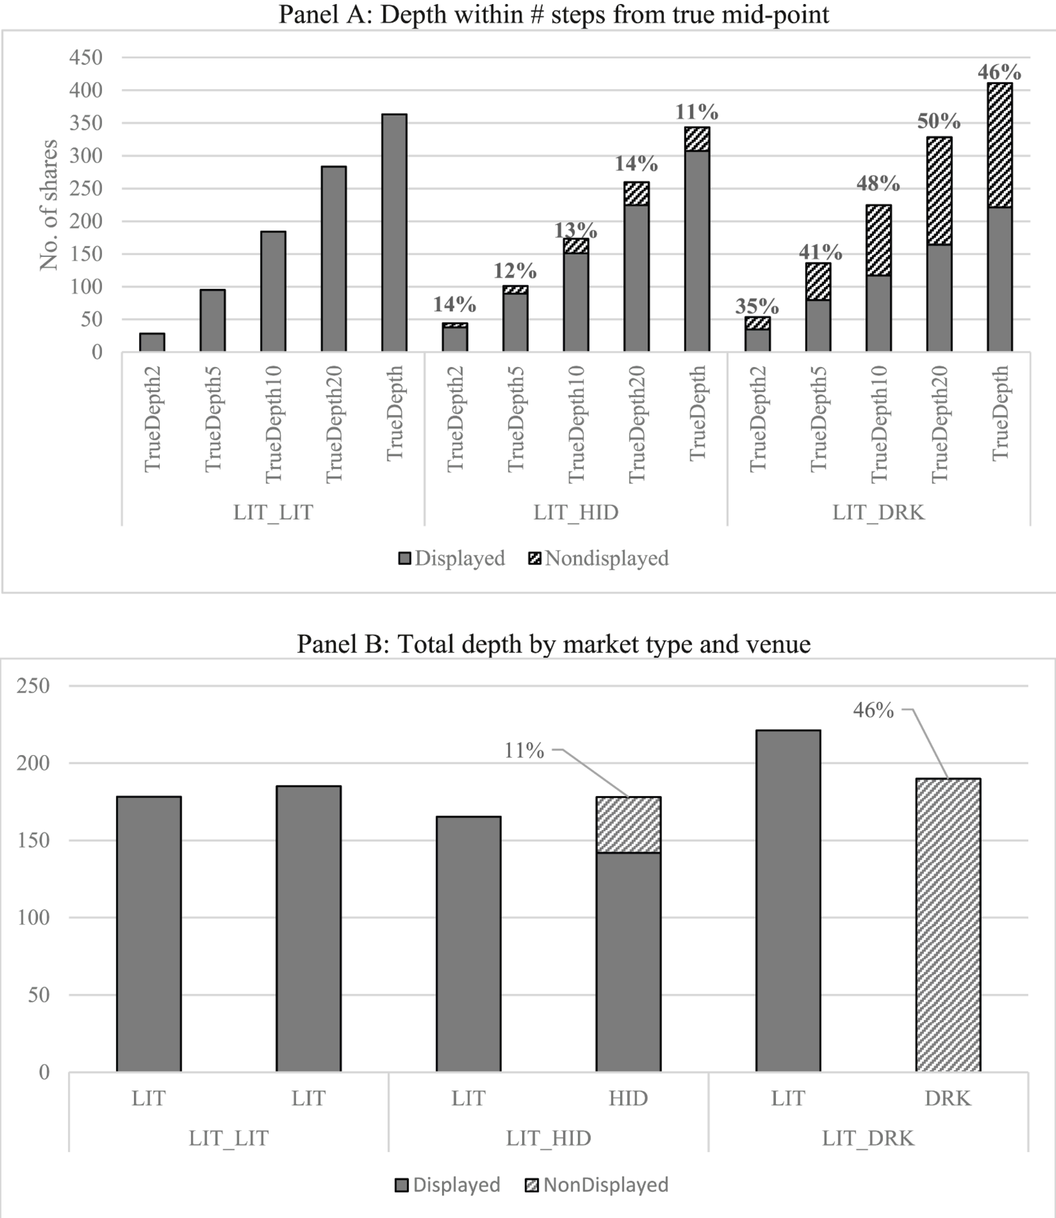

Fig. 3. Depth for each market and the respective venues. Panel A presents depth averaged across cohorts and intervals for each market type, and at various price levels. *TrueDepth2*, *TrueDepth5*, *TrueDepth10*, and *TrueDepth20* are the total number of shares at 2, 5, 10 or 20 steps away from the true midpoint, respectively. The total depth (*TrueDepth*) comprises both displayed and nondisplayed shares at all price levels. The shaded portions represent the nondisplayed depth. Observations within each 30-s interval are time-weighted to obtain an interval-level observation. The interval observations are then averaged for each security. In the figures, we plot the average across 48 observations (i.e., 12 cohorts of 4 securities for each market type). Panel B presents the total depth measures averaged across cohorts and intervals for each market and the respective venues.

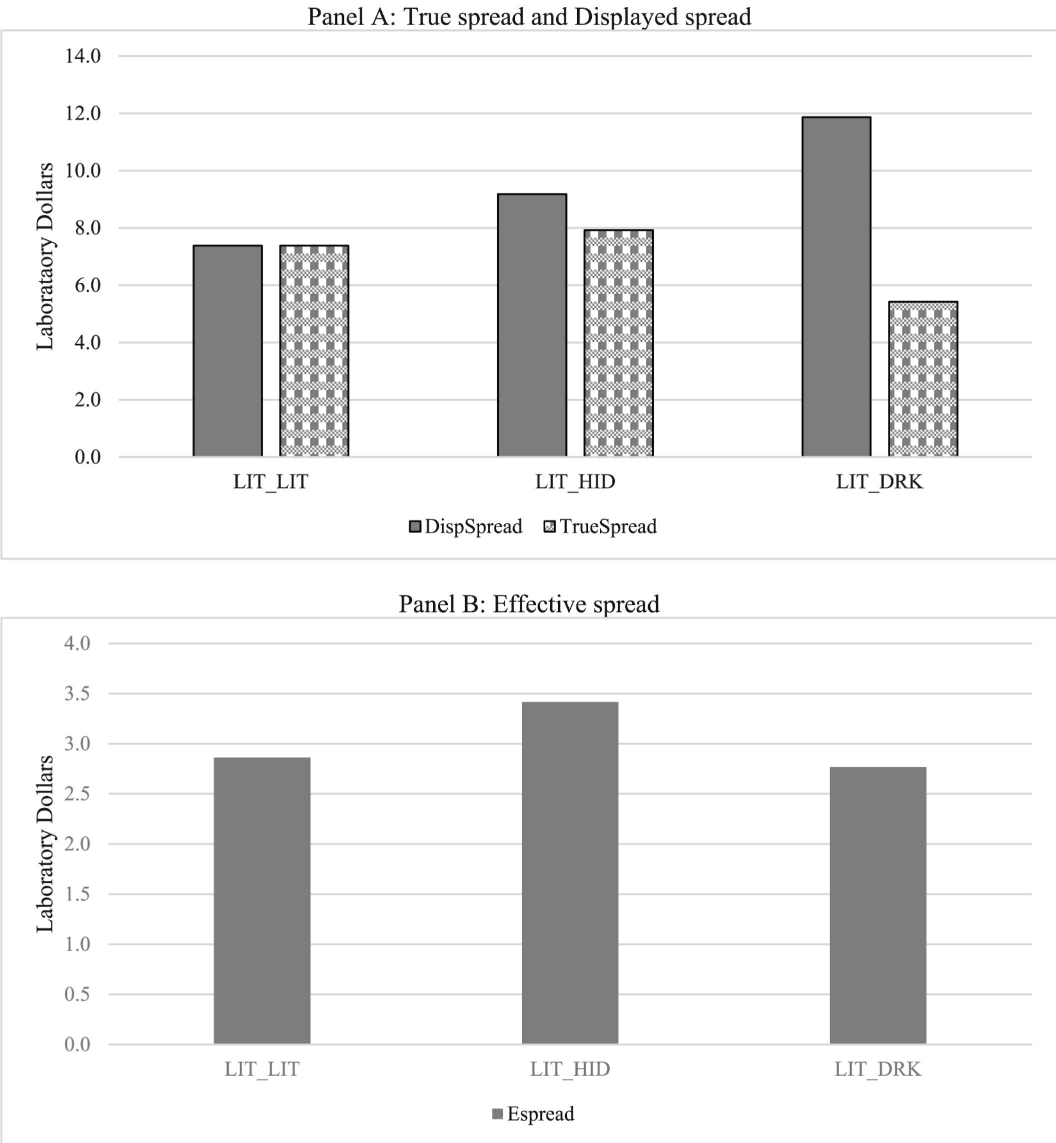

Fig. 4. Cross-sectional comparison of spread measures. This figure presents the average value of the spread measures for each market type. In Panel A, *TrueSpread* is the difference between the best "true" ask and best "true" bid, where "true" best bid and best ask prices is determined by examining both displayed and nondisplayed shares. *DispSpread* considers only the displayed best bid and best ask prices. Both measures are averaged across cohorts and intervals for each market type. In Panel B, Effective spread, *Espread,* is defined as the traded price minus the midpoint of bid-ask spread when trade occurs for buyer-initiated trade and the midpoint minus the traded price for seller-initiated trade. The measure is averaged across cohorts for each market type.

Table 7

Transaction costs for informed traders and liquidity traders by market type. This table presents the average effective spread (*Espread*), realized spread (*Rspread*) and price impact component (*PImpact*) at the market level, by whether the trader is trading as a liquidity demander (Panel A) or supplier (Panel B), and by informational roles (i.e., Informed and Liquidity). Panel C presents the same measures without considering whether the trader is trading as liquidity demander or supplier. Effective spread (*Espread*) is defined as the traded price minus the midpoint of bid-ask spread when trade occurs for buyer-initiated trade and the midpoint minus the traded price for seller-initiated trade. *Espread* is decomposed into Realized spread (*Rspread*) and the price impact component (*PImpact*). *Rspread* is defined as the difference between traded price and true value. *PImpact* is the difference between true value and the midpoint of bid-ask spread when the trade occurs. The measures presented are averaged across 48 observations (i.e., 12 cohorts of 4 securities for each market type). The repeated measures ANOVA tests are based on 192 market-trader type level observations.

		LIT_LIT (1)	LIT_HID (2)	LIT_DRK (3)	*t*-test Informed against Liquidity (*p*-value)	ANOVA test (*p*-value)		
						(1) = (2)	(1) = (3)	(2) = (3)
Panel A: Liquidity demander								
Espread	Informed	2.57	2.84	2.41	0.20	0.42	0.66	0.23
	Liquidity	3.09	3.68	2.78				
Rspread	Informed	−3.65	−1.28	−0.84	0.00	0.04	0.40	0.22
	Liquidity	3.15	4.56	1.67				
PImpact	Informed	6.22	4.13	3.25	0.00	0.12	0.32	0.64
	Liquidity	−0.06	−0.87	1.11				
Panel B: Liquidity supplier								
Espread	Informed	−3.62	−3.96	−3.61	0.01	0.37	0.83	0.60
	Liquidity	−2.36	−2.88	−2.60				
Rspread	Informed	−3.57	−4.34	−2.57	0.00	0.19	0.56	0.49
	Liquidity	3.09	1.74	1.08				
PImpact	Informed	−0.05	0.37	−1.03	0.00	0.45	0.67	0.78
	Liquidity	−5.45	−4.62	−3.69				
Panel C: Overall								
Espread	Informed	−1.04	−0.99	−0.81	0.00	0.78	0.16	0.31
	Liquidity	0.47	0.33	−0.20				
Rspread	Informed	−3.94	−2.68	−2.31	0.00	0.16	0.83	0.19
	Liquidity	3.27	2.82	1.76				
PImpact	Informed	2.90	1.69	1.50	0.00	0.09	0.26	0.50
	Liquidity	−2.80	−2.49	−1.96				

spread (*DispSpread*) narrow over the six trading intervals. Similar to the depth measures, the extent of the decline for the spread measures differs across market types.[25]

With regard to market quality, we posit that total depth is larger and the displayed depth is smaller in a market with pre-trade opacity as compared to one without. Panel A in Fig. 3 shows the average total depth in the LIT_DRK market (411 shares) is the highest among the different market types and is statistically different from that in the LIT_HID market (343 shares, two-market p-value = 0.03). However, the average total depth in the LIT_HID market is not statistically different than that in the LIT_LIT market. The displayed depth in the LIT_DRK market is the lowest (221 shares) and is statistically different from that in the LIT_HID market (307 shares, two-market p-value = 0.00) and LIT_LIT market (363 shares, two-market p-value = 0.00). Overall, we find support for our predictions of higher total depth and lower displayed depth in the dark market compared to markets with less opacity.

Panel A of Fig. 3 also presents the average total depth at two, five, ten, and twenty price increments away from the true midpoint (computed using the best ask and best bids of all orders), denoted as *TrueDepth2*, *TrueDepth5*, *TrueDepth10*, and *TrueDepth20*, respectively. When compared across the three markets, the total depth at twenty price steps is significantly larger in the LIT_DRK market compared to the LIT_HID and LIT_LIT markets (three-market p-value = 0.05, two-market p-value = 0.02). The displayed depths at ten and twenty price steps (*DispDepth10* and *DispDepth20*) are both significantly smaller in the LIT_DRK market than the other two markets. Furthermore, when comparing the LIT_DRK market against the LIT_LIT market, the total depth at up to the first five price steps, i.e., *TrueDepth2*, *TrueDepth5*, are significantly larger (p-values = 0.03 and 0.07 respectively), although the total depth for the two markets are not different. This suggests that traders exploit the darkness in the limit order book to hide their trading intentions and tend to be more aggressive in liquidity provision when one of the two venues goes dark. To provide insights on the liquidity provision at the different venues, Panel B in Fig. 3 presents the average total depth at the venue level. The total depth in the LIT and HID venues for the LIT_HID market are not statistically different. However, the total depth is statistically higher in the LIT venue than the DRK venue of the LIT_DRK market (p-value = 0.01).

In the Introduction, we posit that opacity is associated with narrower true spread that considers hidden liquidity and wider displayed spread. Panel A of Fig. 4 shows the mean value of *TrueSpread*, averaged across cohorts and intervals, for the LIT_DRK market is 5.42 laboratory dollars and is significantly lower than those for the LIT_LIT and LIT_HID markets (7.38 and 7.92 with two-market p-

[25] In untabulated results, we examine whether the bid-ask spread is more constrained in particular market types. The bid-ask spread is not often constrained by the tick size across the three market types. For instance, displayed spreads in the LIT_LIT market are constrained in less than 16% of the 30-sec intervals. The spread is most constrained in the last 30-sec interval in the LIT venue of the LIT_DRK market. This is also where we observe an increase in the use of market and marketable limit orders. As tick size in the LIT venue is likely to affect the use of the DRK venue, future research could examine the effects of tick size and transparency in fragmented markets.

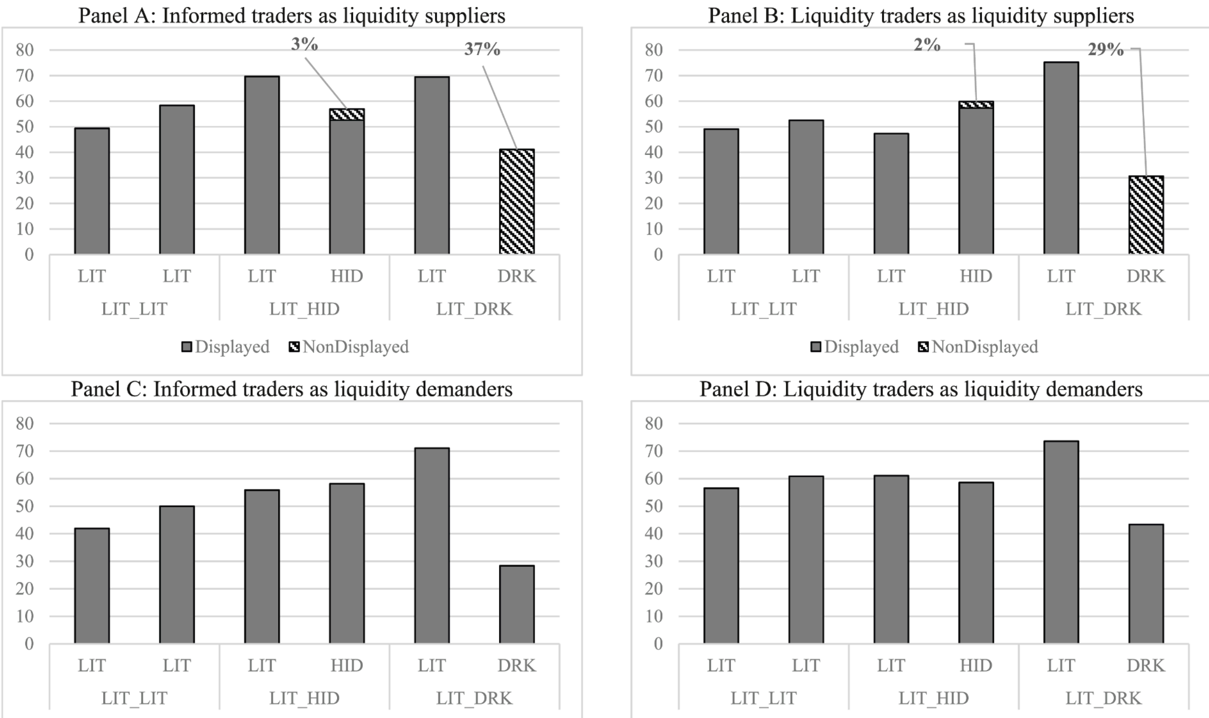

Fig. 5. Trade volume by market type and informational role. This figure presents the trade volume in each venue and market type by whether the trader is trading as a liquidity demander or supplier, and by informational roles (i.e., Informed and Liquidity). In Panels A and B, we plot the average number of shares traded via limit orders by informed and liquidity traders, respectively. In Panels C and D, we plot the average number of shares traded via market orders by informed and liquidity traders, respectively.

values of 0.00 and 0.03, respectively). Conversely, *DispSpread* in the LIT_DRK (11.86 laboratory dollars) is significantly higher than those in the LIT_LIT and the LIT_HID markets (7.38 and 9.18 respectively with two-market p-values of 0.00 and 0.09 respectively). The findings indicate traders use the dark venue to hide their orders between the best displayed bid and ask prices, corresponding to Panel A of Fig. 3. This is consistent with the underlying economic mechanism argued by Boulatov and George (2013). That is, traders cannot condition their order submission decision on the market status in a dark market, hence informed traders tend to be more aggressive in their liquidity provision and cause the true spread to narrow.

The lack of statistical difference between the market liquidity measures for LIT_LIT and LIT_HID suggests the attraction of non-displayed orders [i.e., the reduced exposure cost discussed in Buti and Rindi (2013), and the reduced informational impact modeled by Moinas (2010)] may be mitigated in the multi-venue market due to the risk of being traded-through. Therefore, traders are not more aggressive in liquidity provision in the LIT_HID market than they would otherwise be in the LIT_LIT market, resulting in no difference in *DispSpread* and *TrueSpread* between the two markets. While orders placed in the DRK venue of the LIT_DRK market may encounter the same risk of being traded-through, all orders being completely dark and not be subject to lower display priority, such as in the HID venue, provides incentive for traders to submit orders to the DRK venue.

3.4. *Effects of pre-trade opacity on realized transaction costs*

In light of our prior findings that the true spread is lower and trade-throughs are higher on the LIT_DRK market compared to the other markets, we next analyze traders' effective transaction costs to see whether better true spreads benefit traders or if the benefits are moderated by the trade-throughs. We compute and compare effective spread (*Espread*) across market type as follows, where $q_t = 1$ if there is a buyer-initiated trade and -1 if a seller-initiated trade; p_t is traded price; and mid_t is midpoint of the best bid and ask prices).

$$Espread_t = q_t(p_t - mid_t) \quad (1)$$

In Panel B of Fig. 4, we plot the effective spread across market type and show transaction cost is not statistically significantly affected by pre-trade opacity. The effective spread is the lowest in the LIT_DRK market and the highest in the LIT_HID market, ranging between 2.77 and 3.42 (three-market p-value = 0.42).

We expand the analysis to examine whether opacity affects transaction costs for different trader types when supplying or demanding liquidity. As the true value of the underlying security (*TrueValue*) is known in the laboratory, we also decompose *Espread* into two components: realized spread, $Rspread_t = q_t(p_t - TrueValue)$, and price impact, $PImpact_t = q_t(TrueValue - mid_t)$. Table 7 reports the spread measures for informed and liquidity traders, and when they trade by demanding liquidity or supplying liquidity.

Table 8
Informational efficiency. This table presents the informational efficiency measures by market type. Informational efficiency is measured by the absolute difference between the true value and the quote midpoint each time a trade occurs. Two measures are considered: *TrueDev* is defined as the average of the deviations of the true value from the "true" quote midpoint, incorporating both displayed and nondisplayed bid and ask prices, within the 30-s interval. *Dev* is defined as the average of the deviations of the true value from the midpoint of the displayed bid and ask prices within the 30-s interval. The measures are averaged within-securities (six 30-s intervals), and then averaged across 48 observations (i.e., 12 cohorts of 4 securities for each market type). The repeated measures ANOVA tests, based on 576 market-interval level observations.

	LIT_LIT (1)	LIT_HID (2)	LIT_DRK (3)	ANOVA test (*p*-value)		
				(1) = (2)	(1) = (3)	(2) = (3)
TrueDev	7.25	7.10	6.90	0.82	0.96	0.79
Dev	7.25	7.38	7.10	0.63	0.79	0.87

Across the three markets, informed traders on average incur lower realized spreads (*Rspread* *p*-value = 0.00) than uninformed traders when demanding liquidity. Between market types, informed traders do not earn significantly lower profits (*Rspread*) in the LIT_DRK market than the other two markets (0.84 per trade vs. 1.28 and 3.65 in the LIT_HID and LIT_LIT markets respectively (untabulated *p*-values = 0.74 and 0.06, respectively). The realized spreads for liquidity traders are no different in the LIT_DRK and LIT_LIT markets but are higher in the LIT_HID market. When considering the transactions where traders supply liquidity, we observe that the *Espread* for both trader types are negative as expected across all market types. We also find, in all market types, informed traders earn larger effective and realized spreads than liquidity traders (*Espread* *p*-value = 0.01 and *Rspread* *p*-value = 0.00). More pertinently, we do not find traders earn different transactions costs when supplying liquidity in the different markets.

The documented transaction costs results contrast with our findings shown in Fig. 4, where *TrueSpread* is lowest in the LIT_DRK market. We note that a trader's overall transaction cost is a weighted sum of venue trade volume and *Espread*. Hence the trader's venue trade volume can significantly affect the transaction cost. Fig. 5 presents the trade volume by market type, informational role, and the order type used (i.e., liquidity demanders vs. suppliers). In Panels A and B, we observe, across the three markets, that liquidity traders are less likely to trade using limit orders than informed traders. In Panels C and D, we observe that liquidity traders demand more liquidity than informed traders. This suggests liquidity traders would have a higher transaction cost as they are more likely to pay rather than earn the adverse-selection component of the spread. We also note that informed traders are less likely to demand liquidity in the DRK venue than in the LIT venue of the LIT_DRK market. It is evident from Fig. 5 that the trading volume in the DRK venue of the LIT_DRK market is significantly lower than the LIT venue. While the spread is lower in the LIT_DRK market (see *TrueSpread* in Panel A of Fig. 4), traders do not benefit from the lower spread as volume is low in the DRK venue. The evidence suggests the cost to traders in checking the best quotes in the DRK venue, such as by pinging the venue with small orders, often proved to be prohibitive in our experimental market.[26]

3.5. Effects of pre-trade opacity on informational efficiency

One advantage for adopting an experimental approach is that the fundamental value of the traded asset is well-defined beforehand, and this enables us to study price efficiency, which is difficult to measure using observational data. We examine the deviations of the "true" quote midpoint (*TrueDev*) and the quote midpoint of the displayed bid and ask prices (*Dev*) from the true value. Table 8 shows the deviation of true value from the quote midpoint of the displayed bid and ask prices and from the true midpoint (based on both displayed and nondisplayed shares) are not statistically different across the three market types. Hence, our results do not lend support to concerns with the adverse effect of dark pools on price discovery.

4. Conclusions

Transparency is an important topic for regulators, practitioners, and academics. We experimentally study pre-trade transparency in fragmented markets and compare how the ability to hide orders in the presence of endogenous and exogenous opacity affects traders' trading strategies, market liquidity, and traders' transaction cost. Our evidence supports the view that opacity encourages liquidity provision, but this benefit is offset by increased trade-throughs in fragmented markets. Our evidence also shows different forms of pre-trade opacity yield somewhat different market outcomes.

When two trading venues are available and one goes dark, traders substitute nondisplayed orders for displayed ones and on average they provide more liquidity than when both venues are lit. The informed traders acting as liquidity suppliers in the dark venue compete

[26] We investigate the use of "pinging" orders by the traders in our experimental setting by examining two trading strategies: (1) front-running, which involves the use of a market order in the DRK venue, followed by a market or limit order on the different side to the market order in the LIT venue, and (2) seeking liquidity, which involves the use of a market in the DRK venue, followed by a market or limit order on the same side to the market order in the LIT venue. We find few instances of such orders and suggest traders may find it difficult to implement such strategies due to the short duration of each game.

in liquidity provision. This results in the "true" bid-ask spread that incorporates both displayed and nondisplayed shares being narrower and the true depth being greater. However, these changes in traders' trading strategies and market quality do not significantly affect traders' overall transaction cost. This is mainly driven by: (1) traders being unable to take advantage of the narrower "true" bid-ask spread due to the greater prevalence of trade-throughs with increased pre-trade opacity; and (2) the higher execution risk due to lower volume and fill rates in the dark venue. These suggest that changes in transparency likely affect different types of investors and traders differently based on their sophistication, their ability to locate the best price, and where they trade.

In the presence of endogenous opacity, traders can choose to hide shares using partially nondisplayed orders. When one venue permits partially nondisplayed orders, only about 10% of depth is hidden and the overall market outcomes do not significantly differ from those in a market with two identical lit venues. The only market quality measure that differs is decreased displayed depth.

In practice, the introduction of dark trading not only introduces darkness in the limit order book, but also increases market fragmentation. By controlling for market fragmentation in our experiments, we provide evidence supporting the view that darkness in the limit order book does not harm market quality nor informational efficiency. In other words, we do not find evidence that opacity worsens price discovery, a concern expressed by market regulators. However, we also do not find evidence that increasing transparency benefits the market. Practical complexities not in our experiment, such as smart-order routing technology and derivative pricing rules, warrant future study.

Acknowledgements

We are grateful to Paolo Pasquariello (the editor), two anonymous referees, Gideon Saar, Maureen O'Hara, Vincent Crawford, Huu Duong, and Arie Gozluklu for their helpful comments and suggestions. This paper also benefitted from comments received at research seminars at Monash University, Deakin University, University of Western Australia, and the West Coast Finance Colloquium (2020, Perth). Financial support for the project was provided by the Accounting and Finance Association of Australia and New Zealand (AFAANZ Research Grant, 2016) and University of Western Australia (BHP Billiton Research Awards, 2016).

Appendix 1. Details of market type and trading platform

Table A

Description of market type. In this table, we summarize the key features of four market types in our experiments. They are: (1) A single lit limit order book market (LIT); (2) Two identical lit limit order book markets (LIT_LIT); (3) Two limit order book markets, one of which allows hidden orders and iceberg orders (LIT_HID); and (4) Two limit order book markets, one of which does not display limit order book status (LIT_DRK).

Market Type	Description of Different Market Type	Major Differences in Participants' Action Options	Sample Screen[b]
LIT	1. A single venue is available for trading. 2. Limit orders (LMT) and market orders (MKT) are permitted.	Order Entry Panel: Venue = &, Type = LMT or MKT*	FigScreen_A
LIT_LIT	1. Two venues are available for trading. 2. Limit orders (LMT) and market orders (MKT) are permitted in both venues.	Order Entry Panel: Venue = !! or %%, Type = LMT or MKT[a]	FigScreen_B
LIT_HID	1. Two venues are available for trading. 2. In addition to limit orders (LMT) and market orders (MKT), iceberg orders (ICE) are permitted in the alternative venue.	Similar to that of LIT_LIT except that in "Order Entry" panel, Column "Type" lists "ICE" as an option additional to "LMT" and "MKT"	FigScreen_C
LIT_DRK	1. Two venues are available for trading. 2. Limit orders (LMT) and market orders (MKT) are permitted in both venues. However, the alternative venue is not visible.	Similar to that of LIT_LIT except that "Alternative Market Depth is hidden"	FigScreen_D

[a] We use the following symbols in our labelling of the venues: &, @, #, $, %, ˆ, ! and *. For a LIT market, we label the venue by one character. For the other market types, we use 2–3 characters to create labels and then randomly assign these labels to the trading venues. The list of venue labels are available upon request.

[b] In the screen layout for the LIT_LIT, LIT_HID, and LIT_DRK markets, the information of one venue is placed above the other. To avoid any effect arising from participants' preference of a venue due to the position of the venue shown on the screen, two screen layouts are used during each game and are randomly assigned to the eight participants. The screen layouts differ only by which venue is listed first in the Market View and Market Dept.

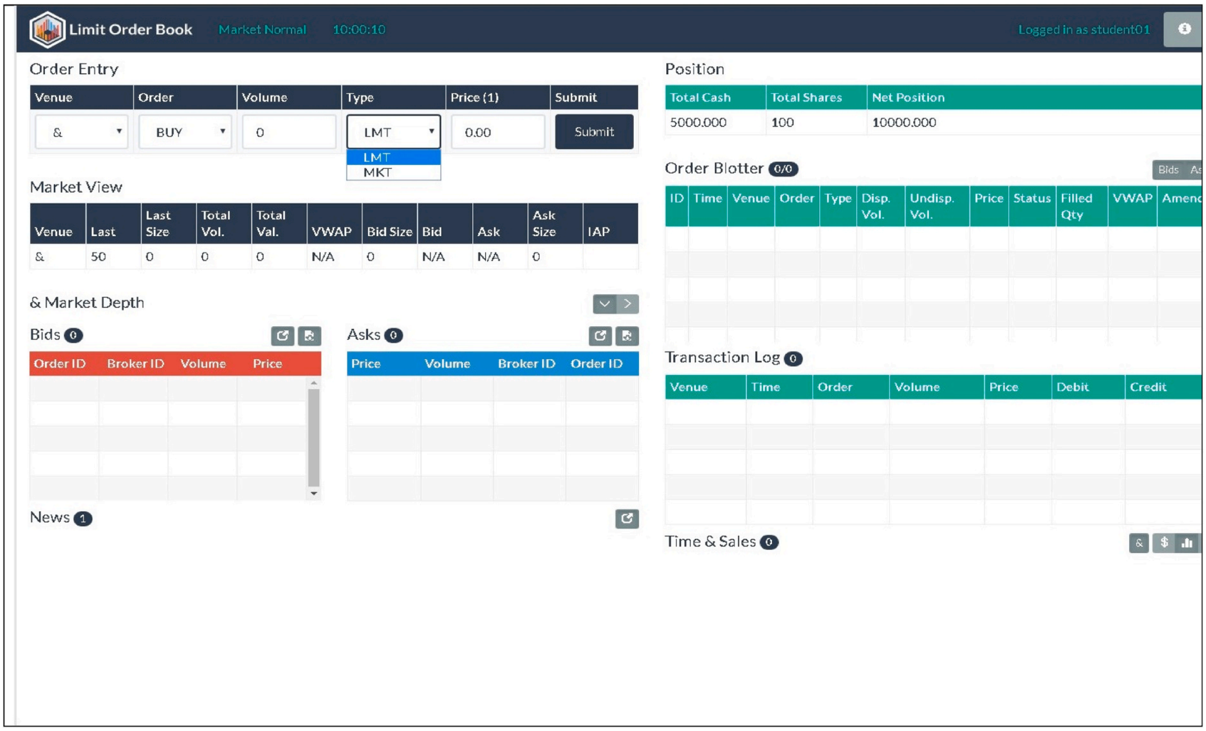

FigScreen_A. Sample screen of a LIT market.

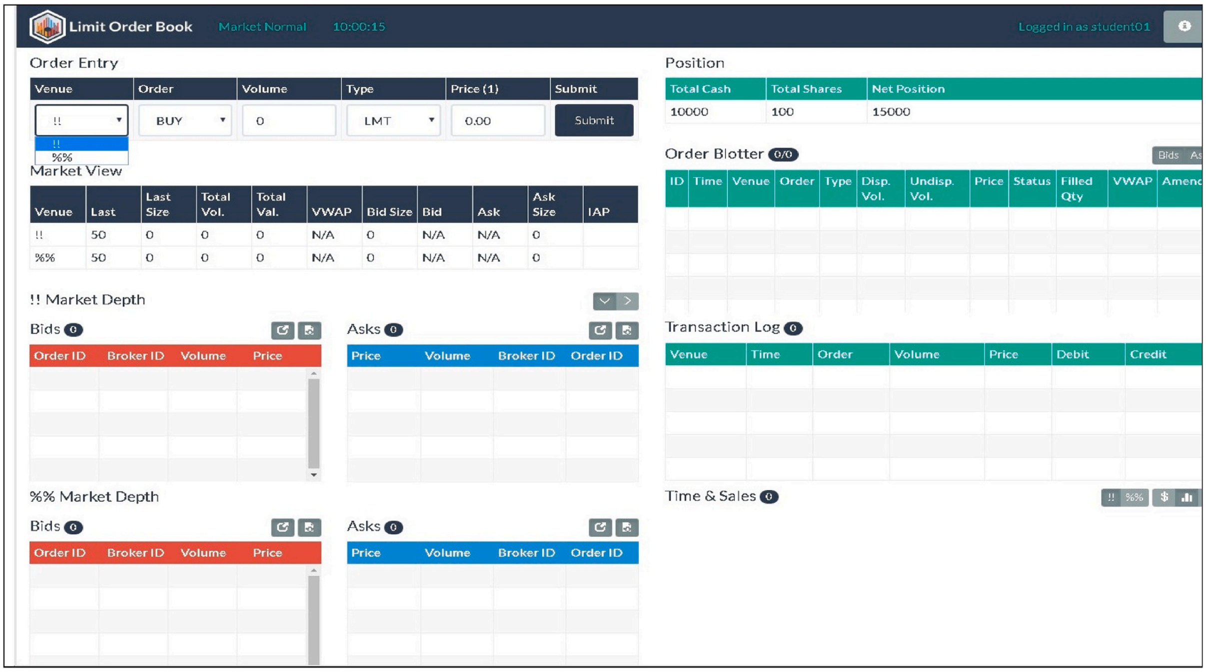

FigScreen_B. Sample screen of a LIT_LIT market.

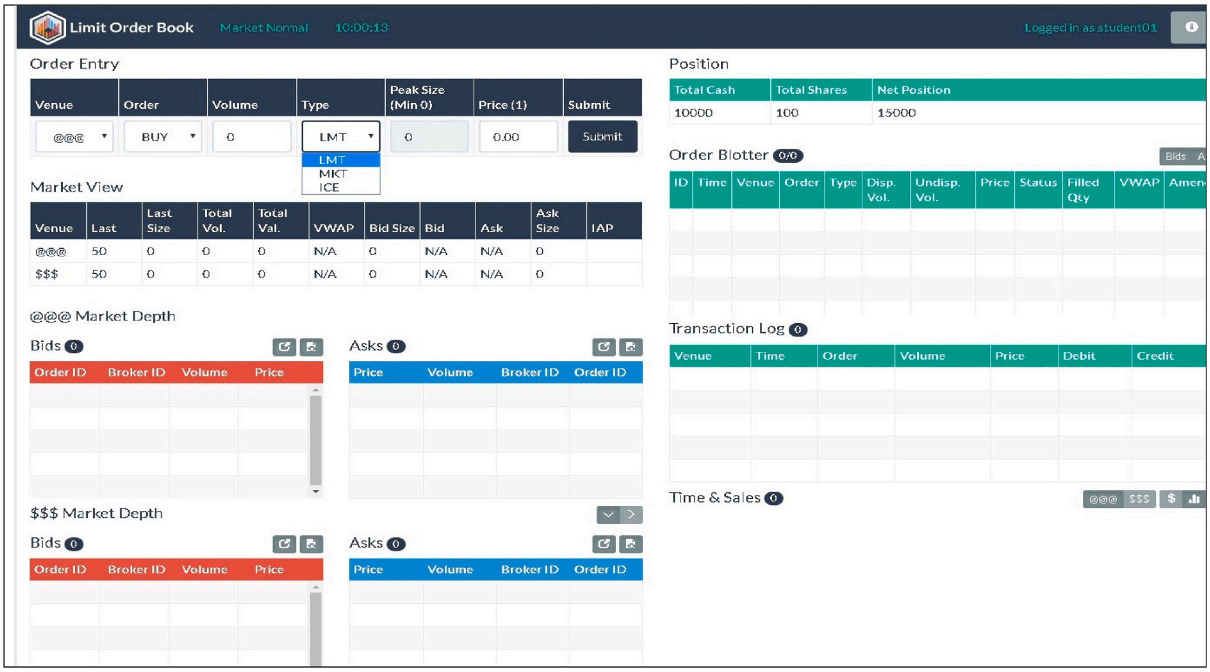

FigScreen_C. Sample screen of a LIT_HID market.

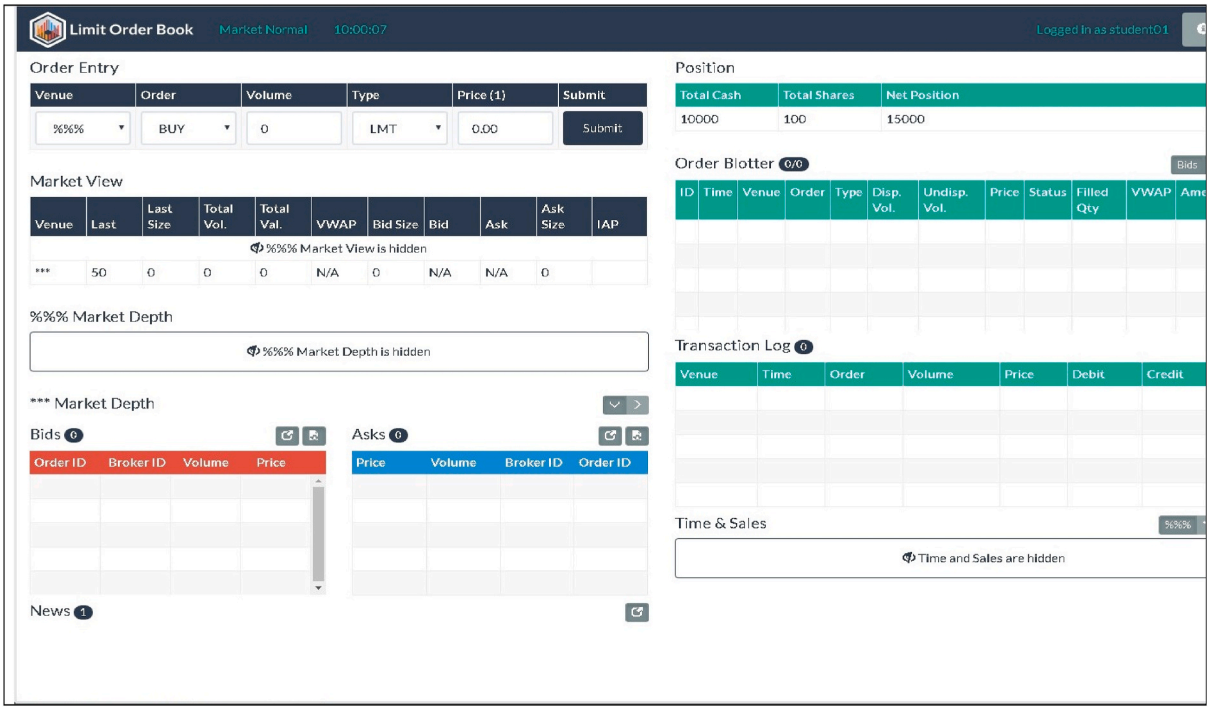

FigScreen_D. Sample screen of a LIT_DRK market.

Appendix B. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.finmar.2022.100732.

References

Anand, A., Samadi, M., Sokobin, J., Venkataraman, K., 2021. Institutional order handling and broker-affiliated trading venues. Rev. Financ. Stud. 34 (7), 3364–3402. https://doi.org/10.1093/rfs/hhab004.

Aquilina, M., Foley, S., O'Neill, P., Thomas, R., 2016. Asymmetries in Dark Pool Reference Prices. FCA Occasional Paper No. 21, Available at: SSRN: https://ssrn.com/abstract=2848120.

Battalio, R., Hatch, B., Jennings, R., 2004. Toward a national market system for U.S. Exchange–Listed equity options. J. Finance 59 (2), 933–962. https://doi.org/10.1111/j.1540-6261.2004.00653.x.

Bayona, A., Dumitrescu, A., Manzano, C., 2021. Information and Optimal Trading Strategies with Dark Pools. Available at: SSRN: https://ssrn.com/abstract=2995956.

Bloomfield, R., O'Hara, M., Saar, G., 2005. The "make or take" decision in an electronic market: evidence on the evolution of liquidity. J. Financ. Econ. 75 (1), 165–199. https://doi.org/10.1016/j.jfineco.2004.07.001.

Bloomfield, R., O'Hara, M., Saar, G., 2015. Hidden liquidity: some New light on dark trading. J. Finance 70 (5), 2227–2274. https://doi.org/10.1111/jofi.12301.

Boulatov, A., George, T.J., 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. Rev. Financ. Stud. 26 (8), 2096–2137. https://doi.org/10.1093/rfs/hht035.

Buti, S., Rindi, B., 2013. Undisclosed orders and optimal submission strategies in a limit order market. J. Financ. Econ. 109 (3), 797–812. https://doi.org/10.1016/j.jfineco.2013.04.002.

Buti, S., Rindi, B., Werner, I.M., 2017. Dark pool trading strategies, market quality and welfare. J. Financ. Econ. 124 (2), 244–265. https://doi.org/10.1016/j.jfineco.2016.02.002.

Brolley, M., 2020. Price improvement and execution risk in lit and dark markets. Manag. Sci. 66 (2), 863–886. https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2018.3204.

Comerton-Forde, C., Putniņš, T.J., 2015. Dark trading and price discovery. J. Financ. Econ. 118 (1), 70–92. https://doi.org/10.1016/j.jfineco.2015.06.013.

Costa-Gomes, M., Crawford, V.P., Broseta, B., 2001. Cognition and behavior in normal-form games: an experimental study. Econometrica 69 (5), 1193–1235.

Foley, S., Putniņš, T.J., 2016. Should we be afraid of the dark? Dark trading and market quality. J. Financ. Econ. 122 (3), 456–481. https://doi.org/10.1016/j.jfineco.2016.08.004.

Foucault, T., Menkveld, A.J., 2008. Competition for order flow and smart order routing systems. J. Finance 63 (1), 119–158. https://doi.org/10.1111/j.1540-6261.2008.01312.x.

Gozluklu, A.E., 2016. Pre-trade transparency and informed trading: experimental evidence on undisclosed orders. J. Financ. Mark. 28, 91–115. https://doi.org/10.1016/j.finmar.2016.03.003. Supplement C.

Harris, L., 2015. Transaction Costs, Trade Throughs, and Riskless Principal Trading in Corporate Bond Markets. Available at: SSRN: https://ssrn.com/abstract=2661801.

Harris, L., Mehta, A., 2020. Riskless Principal Trades in Corporate Bond Markets. Available at: SSRN: https://ssrn.com/abstract=3681652.

Hendershott, T., Jones, C.M., 2005. Trade-through prohibitions and market quality. J. Financ. Mark. 8 (1), 1–23. https://doi.org/10.1016/j.finmar.2004.09.001.

Kovaleva, P., Iori, G., 2015. The impact of reduced pre-trade transparency regimes on market quality. J. Econ. Dynam. Control 57, 145–162. https://doi.org/10.1016/j.jedc.2015.05.011. Supplement C.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53 (6), 1315–1335.

Menkveld, A.J., Yushen, B.Z., Zhu, H., 2017. Shades of darkness: a pecking order of trading venues. J. Financ. Econ. 124 (3), 503–534. https://doi.org/10.1016/j.jfineco.2017.03.004.

Moinas, S., 2010. Hidden Limit Orders and Liquidity in Order Driven Markets. Available at: http://publications.ut-capitole.fr/id/eprint/3350.

Peterson, M., Sirri, E., 2002. Order submission strategy and the curious case of marketable limit orders. J. Financ. Quant. Anal. 37 (2), 221–241. https://doi.org/10.2307/3595004.

Securities and Exchange Commission, 2013. Staff Report on Equity Market Structure Literature Review Part I: Market Fragmentation. Available at: https://www.sec.gov/files/marketstructure/research/fragmentation-lit-review-100713.pdf.

Securities and Exchange Commission, 2020. Staff Report on Algorithmic Trading in U.S. Capital Markets. Available at: https://www.sec.gov/files/Algo_Trading_Report_2020.pdf.

White, M.J., 2014. Enhancing Our Equity Market Structure. Available at: https://www.sec.gov/news/speech/2014-spch060514mjw.

Zhu, H., 2014. Do dark pools harm price discovery? Rev. Financ. Stud. 27 (3), 747–789. https://doi.org/10.1093/rfs/hht078.