On the Effects of Continuous Trading*

Ivan Indriawan† **Roberto Pascual**‡ **Andriy Shkilko**§

November 14, 2021

Abstract. The continuous limit order book is a prominent design feature of modern securities markets. Theoretical models show that this feature is prone to generating adverse selection and recommend switching to batch auctions as a superior alternative. We examine an opposite move, whereby a modern stock exchange switches from auctions to continuous trading. Consistent with theory, adverse selection substantially increases. The increase is partly offset by reductions in other trading cost components. Trading volume increases, likely driven by latency arbitrage. Our results suggest that market design optimization is an intricate balancing act and help explain the current dominance of the continuous design.

Key words: liquidity, price efficiency, continuous trading, batch auctions

JEL: G14; G15

*We thank Michael Brolley, Eric Budish, Benjamin Connault, Julio Crego, Björn Hagstömer, Pankaj Jain, Kai Li, Jane Liu, Sophie Moinas, Pamela Moulton, Esen Onur, Roberto Riccó, John Shim, Konstantin Sokolov, Christian Westheide, Marius Zoican, and seminar/session participants at the Finance Forum of the Spanish Finance Association, Financial Markets and Corporate Governance Conference, Mid-Atlantic Research Conference in Finance, Midwest Finance Association Conference, Plato Market Innovator Conference, Stern/Salomon Microstructure Conference, Macquarie University, and Wilfrid Laurier University for valuable comments. Pascual acknowledges the Ministerio de Ciencia, Innovación y Universidades (MCIU), the Agencia Estatal de Investigación (AEI), and the European Regional Development Funds (ERDF) for their support of project ECO2017-86903-P and additional financial support of Generalitat Valenciana Grant Prometeo/2017/158. Shkilko acknowledges support of the Social Sciences and Humanities Council of Canada and the Canada Research Chairs program. The most recent version of this paper may be downloaded from: https://bit.ly/2F8mSvO.

†Auckland University of Technology, New Zealand, e-mail: ivan.indriawan@aut.ac.nz

‡University of the Balearic Islands, Spain, e-mail: rpascual@uib.es

§Wilfrid Laurier University, Canada, e-mail: ashkilko@wlu.ca

1. Introduction

Many modern markets are organized as continuous limit order books. In this design, market participants submit messages in continuous time, and exchange matching engines process the messages one by one in order of receipt. Theoretical literature argues that this design may reduce the ability of liquidity providers to reprice stale quotes before they are picked off by latency arbitrageurs (e.g., Budish, Cramton, and Shim (2015)). The adverse selection cost of market making increases as a result, and the increase is relayed to liquidity consumers through greater trading costs. As a remedy, the literature proposes replacing continuous trading with frequent batch auctions, in which orders accumulate for a brief period of time before being matched against each other, providing market makers with an opportunity to change stale quotes.

Empirical studies have not yet directly examined these theoretical predictions, largely because switches between the two market designs are rare. We fill this gap by studying a recent move by the Taiwan Stock Exchange (TWSE) to transition all of its activity from batch auctions to continuous trading. Our main results support the theoretical predictions. In a difference-in-differences (DID) setup, we find that continuous trading is associated with significantly greater adverse selection and greater trading costs.

As latency arbitrageurs pick off stale quotes, they generate trades. Shkilko and Sokolov (2020) and Aquilina, Budish, and O'Neill (2021) show that these trades add up to a substantial share of trading volume. Since continuous trading stimulates latency arbitrage, total trading volume may increase after the TWSE transitions to the new regime. Alternatively however, the above-mentioned increase in trading costs may cause some market participants to scale down their activity, reducing total volume. When we take these possibilities to the data, we observe that trading volume increases overall. This finding highlights an important tension in modern markets. On the one hand, exchange revenues are volume-dependent, so exchange operators prefer market designs that maximize volume. On the other hand, such designs may not benefit all market participants, especially those who seek liquidity. Consistent with Budish, Lee, and Shim (2021), private-market incentives may therefore be insufficient for welfare maximization.

Adverse selection is an important, yet not the only, concern of a market maker. Among her other considerations are inventory costs, fixed costs, and profits. We are not aware of a theoretical model that examines the relation between these three considerations and continuous trading, so in a later section we discuss a range of possibilities. Empirical studies usually aggregate the three into a composite metric called the realized spread. The data show that this metric decreases after the TWSE moves to continuous trading. The magnitude of the decrease is smaller than that of the increase in adverse selection, and therefore trading costs increase overall. This result provides some background for future theoretical work and also cautions that optimizing market design may be a delicate balancing act that involves several offsetting factors.

The 21st century has witnessed significant changes in the structure of financial markets. Exchanges have largely automated the trading process (Hendershott, Jones, and Menkveld (2011), Hendershott and Moulton (2011)) and considerably improved matching engine connectivity and execution speeds (Conrad, Wahal, and Xiang (2015), Brogaard, Hagströmer, Nordén, and Riordan (2015)). Market participants responded to these changes by adopting the latest technology in a speed race to the exchange engines and between markets (Shkilko and Sokolov (2020)). One market structure feature that has however remained largely unchanged during this time is the continuous limit order book. In it, orders are fed into the exchange engine one at a time on a first come, first served basis. In the event of two orders arriving simultaneously, chance determines which order is processed first.

Budish, Cramton, and Shim (2015) question this design due to its ability to intensify adverse selection through the latency arbitrage channel. To understand their reasoning, it helps to think of a group of *N* market participants, who have identical speeds, all reacting to the same information. All *N* participants may act both as market makers and liquidity takers (snipers). In the former role, they rush to change their posted quotes in response to news, while in the latter role, they attempt to pick off the stale quotes of others. Even though everyone's speeds are the same, chance dictates that one order will be processed by the exchange engine first. Given that there are $N-1$ snipers for each stale quote, the odds of being adversely selected, $(N-1)/N$, are not in favour of the market maker. In the meantime, a batch auction that accumulates orders for a period of time before matching them gives the market maker sufficient time to revise her stale quote before it is picked off. As long as the auctions are not ultra-frequent, as argued by Haas, Khapko, and Zoican (2020), she can do so even if the other traders are a little faster.

Latency arbitrage opportunities have many origins, ranging from publicly observable information conduits specific to the stock to outside sources such as correlated assets and various public announcements. For an example of a stock-specific conduit, recall that informed market participants regularly use limit orders (O'Hara (2015), Brogaard, Hendershott, and Riordan (2019)), and therefore the stock's own limit order book imbalance often serves as an early signal of future price movements. Recent theoretical models, i.e., Riccó, Rindi, and Seppi (2020) and Bhattacharya and Saar (2020), expand the set of stock-specific conduits to price history, the state of liquidity, order arrival frequency, and order submission timing. Brogaard, Hendershott, and Riordan (2014) and Kwan, Philip, and Shkilko (2021) show that sophisticated traders routinely use signals derived from the conduits to facilitate informed trading, thus generating adverse selection. To provide a brief illustration, let us assume that the book imbalance becomes substantially positive. Since this positive signal is visible by many market participants, it is likely to trigger latency races, with market makers trying to reprice their quotes, and snipers attempting to pick off these quotes. In a continuous market, these races are likely to result in adverse selection.

It is possible that such races are more common in high-volume stocks. First, trading frequency should correlate with the number of limit order book updates as well as with changes in the other stock-specific conduits. Second, when it comes to outside sources, Shkilko and Sokolov (2020) show that latency arbitrage tends to be more pronounced between correlated assets that trade frequently. With this in mind, we examine the results in the cross-section. To do so, we divide the sample into trading volume terciles and re-estimate the main regression models for each tercile. Consistent with expectations, the data show that adverse selection increases in all terciles, but the increase is substantially more pronounced in the frequently-traded stocks. Notably, the relatively small adverse selection increase in the lower-volume stocks is virtually entirely offset by the reduction in the realized spreads, for a zero net effect on trading costs. As such, the cross-section corroborates our earlier suggestion that optimizing market design may be an intricate balancing act.

It is possible that by picking off the stale quotes, latency arbitrage brings prices closer to their true values thereby improving price efficiency. It is also possible however that liquidity providers were already maintaining efficiency at the optimal level during the discrete regime by promptly adjusting their quotes. Price efficiency is an important aspect of market quality, so we take these possibilities to the data. Using several conventional metrics, we find that price efficiency effects of continuous trading are mixed; some metrics improve, some deteriorate, and others remain unchanged. Overall, continuous trading does not appear to be obviously superior to batch auctions in its ability to deliver efficient prices.

The TWSE is one of the world's 20 largest stock exchanges. Ranked by the U.S. dollar trading volume, it is comparable (ranked 15$^{\text{th}}$) to such markets as the Toronto Stock Exchange (13$^{\text{th}}$) and the Australian Securities Exchange (20$^{\text{th}}$). Until recently, the TWSE was the only large market that used batch auctions as the primary method of matching buyers and sellers. The auctions were relatively frequent, occurring every five seconds, yet recently the exchange joined its industry peers in offering continuous market access. Its new continuous trading platform launched on March 23, 2020. It is important to acknowledge that the TWSE switched to continuous trading at the onset of the COVID-19 pandemic, and therefore we must be careful with inferences. To allay concerns with confounding effects, in all analyses we rely on a two-pronged approach.

First, we use a DID setup with a control sample of stocks trading on the Korean Stock Exchange (KRX). The similarities in pandemic responses undertaken by Taiwan and South Korea allow us to cautiously assert that the DID analysis mitigates the confounding effects of the pandemic onset. Second, we use several event window lengths, some of which are considerably removed from the pandemic onset period, to enhance robustness. The results are preserved regardless of the event window length or event window proximity to March 23. Taken together, these analyses give us sufficient confidence that the findings are attributable to the switch to continuous trading rather than the pandemic. We note that due to the one-event nature of the TWSE switch, a DID analysis would have been prudent even in the absence of the pandemic. For such an analysis, the geographic proximity of the two markets and their similar sizes would have made the KRX a sensible source of controls.

To date, the proposal to discretize trading has not gained much traction in the exchange industry. Only one U.S. market operator, Cboe Global Markets, has obtained approval from the Securities and Exchange Commission to implement batch auctions on one of its smaller equity exchanges, BYX.[1] Our results help explain the general reluctance of the industry to change the status quo. Recall that continuous trading comes with an increase in trading volume, an important revenue driver for modern exchanges. In an industry characterized by high fixed costs, willfully reducing a revenue source is generally inconsistent with profit maximization. Notably in this regard, Cboe plans to preserve the BYX continuous book and to run the auctions alongside it.

When the BYX periodic auctions launch, it may be of interest to compare their market quality effects to those obtained for the TWSE. We however caution that the multi-market environment that characterizes U.S. equity trading may not be ideally suitable for such a comparison. Adding a batch auction market to the existing continuous markets may result in a clientele migration and therefore confound market quality inferences. A similar concern may accompany analyses of periodic auctions in Europe, which too operate in a fragmented environment (Johann, Putniņš, Sagade, and Westheide (2019)). In the meantime, the TWSE transition to continuous trading

[1] "Cboe Receives Regulatory Approval to Launch Periodic Auctions for U.S. Equities Trading," March 29, 2021 (http:/bit.ly/2O6wI61).

occurs in a highly consolidated marketplace and therefore suits our purposes rather well.

Taken together, our contribution to the literature is four-fold. First, we empirically confirm theoretical predictions that continuous trading is associated with greater adverse selection. Second, we show that the trading cost components unrelated to adverse selection decrease in the continuous regime. In the frequently-traded stocks, the adverse selection effect dominates, and trading costs increase. In the stocks that trade less frequently, the two effects cancel each other out. Third, we show that continuous trading benefits exchanges by increasing total trading volume – a major source of exchange revenue. As such, our results are consistent with the theoretical predictions that private incentives of exchange operators may not be sufficient to seek alternatives to the continuous design. Our fourth contribution is to adapt traditional trading cost metrics to the periodic auction environment. We describe this adaptation in the next section. In the Internet Appendix, we further expand the discussion of related literature.

2. Data and metrics

2.1 Sample

We collect intraday quote and trade data from the Refinitiv Tick History database, the successor to the Thomson Reuters Tick History database. The sample consists of 100 TWSE stocks with the largest market capitalization. The sample period is nine months, from November 2019 through July 2020. To establish a baseline, Table 1 reports summary statistics computed in the first three months of the sample period, that is November-December 2019 and January 2020. We use this three-month period as the main pre-event window in all subsequent analyses.

The average sample stock has a market capitalization of 283 billion New Taiwan dollars (NTD), share price of NTD 174.18, and daily volume of about 11.55 million shares. The average trade size is 7,313 shares, and volatility (computed as the difference between the highest and lowest daily midquotes scaled by the highest midquote) is 0.016. The sample covers a broad cross section, with market capitalizations ranging between NTD 54 billion and 476 billion (respectively, in the 10^{th} and 90^{th} percentiles), prices ranging between NTD 14.40 and 333.13, and daily volumes – between 0.68 and 25.10 million shares.

[Table 1]

2.2 Liquidity metrics in the continuous regime

Liquidity analyses in continuous markets are quite routine. Meanwhile, liquidity in auction environments is examined less often, and comparisons between continuous and auction regimes are even less common. As such, we set out to carefully explain our measurement approach. To establish a baseline, we begin by describing conventional liquidity metrics for continuous trading and follow with a discussion of comparable metrics for auctions.

Upon switching to continuous trading, the TWSE begins reporting trade and quote data in a format similar to that of the Trade and Quote Database often used to examine liquidity in the U.S. The data contain intraday activity at the top of the limit order book including trades, ask and bid quotes, and quoted depths time-stamped to the nearest millisecond. We bunch trade records that have the same time stamp, trade direction, and price into one trade, as such records typically reflect a trade initiated by one market participant that executes against several standing limit orders. As is common, we also omit the market opening and closing periods for each trading day.

To assess displayed liquidity, we estimate the *quoted spread* as the difference between the best offer and the best bid. To measure the number of shares available at displayed prices, we compute *quoted depth* as the average of the best quote sizes. To assess trading costs incurred by liquidity demanders, we compute the *effective spread* as twice the signed difference between the traded price and the midquote at the time of the trade. To measure the levels of adverse selection, we compute the *price impact* as twice the signed difference between the midquote at the time of the trade and the midquote 15 seconds after the trade. O'Hara (2015) suggests that in modern fast markets, 15 seconds is an appropriate horizon to use for such calculations. In the Internet Appendix, we show that the results are robust to computing price impacts at several alternative horizons. Finally, we compute the *realized spread* as the difference between the effective spread and price impact. The realized spread is a composite metric that is often used to gauge (i) the costs of market making that are unrelated to adverse selection and (ii) market maker profits (Hendershott, Jones, and Menkveld (2011), Brogaard, Hagströmer, Nordén, and Riordan (2015)). Because of the composite nature of the metric, its interpretation is somewhat nuanced, and the upcoming discussions carefully take these nuances into account.

We drop instances when the best quotes are locked or crossed, that is when the quoted spread is zero or negative. We also drop block trades, those in the 99$^{\text{th}}$ percentile of trade sizes, as such trades are generally privately negotiated. To sign trades, we rely on the Lee and Ready (1991) algorithm. Chakrabarty, Pascual, and Shkilko (2015) show that this algorithm performs well in modern markets. Finally, we scale all above-mentioned liquidity metrics by the corresponding midquotes.

2.3 Liquidity metrics in the auction regime

To describe the data available to us during the auction regime, we begin with a depiction of the TWSE auction. During the first half of the sample period, the TWSE uses a conventional auction format similar to that discussed by Budish, Cramton, and Shim (2014). The auction aims to bring security buyers and sellers together at the same time and place and compare their demand and supply curves. If the curves intersect, the auction succeeds resulting in a trade. If the curves do not intersect, the auction does not succeed, and no trade takes place.

The auction process consists of two stages: (i) an accumulation stage lasting approximately five seconds, during which orders are being submitted to the exchange engine, and (ii) a much shorter allocation stage, during which orders execute. The two stages together take five seconds. A new accumulation stage begins immediately after the previous allocation stage. The TWSE information policy closely resembles that recommended by Budish, Cramton, and Shim (2015). Namely, after each allocation the exchange publicly reports total volume matched in the auction and the aggregate size of unexecuted buy and sell orders, for five price levels each.

Panel A of Figure 1 contains an example of a successful auction. The buyers seek to purchase a total of 150 shares, of which 20 are sought at NTD 10.00, 50 at NTD 9.99, and so on. Meanwhile, the sellers seek to sell 50 shares at NTD 10.00, another 50 at NTD 10.01, and so on. Supply and demand cross at NTD 10.00, and the auction succeeds for 20 shares. In turn, Panel B contains an example of an unsuccessful auction. This time, the buyers are unwilling to pay more than NTD 9.99, and the sellers are unwilling to accept less than NTD 10.00. No trade takes place.

[Figure 1]

For auctions that are successful, the TWSE data report three items: (i) the number of shares traded (that is, 20 shares in Panel A), (ii) the traded price (NTD 10.00), and (iii) the number of shares that remain unexecuted at the best prices (respectively, 30 shares for sale at NTD 10.00 and 50 shares for purchase at NTD 9.99). For unsuccessful auctions, the data report only the item (iii). For the unsuccessful auction in Panel B, the data will show 50 shares for sale at NTD 10.00 and 70 shares for purchase at NTD 9.99.

Although these auction data do not match continuous data perfectly, they allow us to draw a cautious yet informative comparison. To explain our approach, we turn to an example from continuous markets that uses supply and demand curves. In Panel A of Figure 2, outstanding limit orders result in NTD 9.99 on the bid and 10.00 on the offer. A newly arriving buyer seeking to purchase 20 shares has a choice to make. First, she may choose to be patient and join the queue of bids (Panel B). Second, she may choose to execute quickly and demand liquidity by crossing the spread (Panel C). If she makes the latter choice, the buyer-initiated trade for 20 shares at NTD

10 will execute, and the limit order book will adjust to the state illustrated in Panel D.

[Figure 2]

In this continuous example, the option to join the queue of bids results in a scenario similar to the unsuccessful auction in Figure 1. Meanwhile, the option to cross the spread is akin to the outcome of a successful auction. Given these similarities, we suggest that the prices of unexecuted buy and sell orders in the auction data may be used as proxies for the best quotes. In both panels of Figure 1, this reasoning would point to a quoted spread of NTD 0.01 and a midquote of NTD 9.995.

Although our approach to the auction spread estimation is somewhat novel, it is rather intuitive, especially if viewed through the prism of patient vs. impatient trading. Auction participants may choose to price their orders more or less aggressively. Less aggressively priced orders are akin to limit orders in continuous markets. They may execute if the opposite side seeks immediacy, otherwise they remain standing. In the meantime, aggressively priced orders resemble marketable orders in continuous markets, as they are sufficiently impatient to cross the gap between the patient participants.

Having discussed our approach to spread and midquote estimation, we move to trade signing. We suggest that the Lee-Ready algorithm, a method that is well-established in continuous markets, may also be used for auction trade signing. To explain, we again rely on an illustration. Table 2 contains a sample of the TWSE auction data. The sample contains seven consecutive auctions, with executions occurring every five seconds. The midquote is stable during the first six auctions at NTD 297.75 and decreases to NTD 297.25 as a result of the seventh auction.

[Table 2]

During auctions 1 through 6, although the midquote is stable, the book is heavy on the ask side, suggestive of a potential for a price decline. Consistent with this possibility, in all but one of these auctions trades execute at the bid prices and are therefore likely seller-initiated. The Lee-Ready algorithm signs them accordingly. Auctions 6 and 7 are noteworthy, because they precipitate a reduction in the midquote. Auction 6 sees a large seller-initiated trade that consumes almost all of bid depth at NTD 297.50, creating conditions for a midquote change. Auction 7 achieves this change.

The seventh auction is of particular interest. After the sixth auction, the book has 21,000 remaining shares on the bid at NTD 297.50. During the seventh auction, 38,000 shares execute at this bid price, suggesting that additional 17,000 shares are added to the bid between auctions 6 and 7. Notably, the sellers in auction 7 wish to execute more than 38,000 shares at NTD 297.50, namely 49,000 shares. The 11,000 shares that cannot find a buyer at NTD 297.50 remain unexecuted and are posted as the new ask quote after the allocation stage. This example suggests that benchmarking against the quotes that result from the previous auction, as it is done by the Lee-Ready algorithm, allows for a rather straightforward signing of trades. Using the quotes remaining after auction 6, namely 297.50 on the bid and 298.00 on the ask, the algorithm concludes that the trade executed in auction 7 was initiated by the sellers.

One remaining issue related to liquidity metrics is aggregation. In continuous markets, researchers usually time-weight quoted spreads and depths when computing daily aggregates. In a discrete auction regime however, time-weighting loses its meaning because spreads and depths are reported once per auction, that is once every five seconds. With this in mind, for better comparability between the auction and continuous regimes, we equal-weight quoted spreads and depths through the entire sample period. The remaining liquidity metrics are volume-weighted.

Panel A of Table 3 reports that the average quoted and effective spreads before the switch to continuous trading are, respectively, 23.05 and 22.38 bps, while price impacts and realized spreads are 4.66 and 17.73 bps. Quoted depth is about 427.73 thousand shares. Again, we observe non-trivial variation in the cross-section, with effective spreads for instance ranging from 11.56 bps in the 10$^{\text{th}}$ percentile to 38.49 bps in the 90$^{\text{th}}$ percentile, and price impacts ranging from 0.54

to 10.37 bps. We note that here and in subsequent analyses we winsorize all continuous variables at 1% to mitigate the outlier effects.

[Table 3]

2.4 Price efficiency metrics

In addition to understanding the effects of continuous trading on liquidity costs, we measure its effects on price efficiency. To measure efficiency, we use two standard metrics: *return autocorrelation* as in Hendershott and Jones (2005) and *price delay* of Hou and Moskowitz (2005). The former metric relies on the notion that, in a frictionless market, prices should be unpredictable, and as such midquote returns should have zero autocorrelation. It is defined as the absolute first order midquote return autocorrelation, and we compute it at several frequencies $s \in \{10\text{s}, 30\text{s}, 60\text{s}, 300\text{s}\}$. Smaller autocorrelation estimates suggest greater efficiency.

The latter metric in turn assumes that efficient prices should instantly incorporate public market information. Accordingly, lagged market returns should have no predictive power for individual stocks returns. To compute this metric, we begin by running the following regression for each stock-day:

$$r_{i,s} = \alpha_i + \beta_i r_{m,s} + \sum_{k=1}^{10} \gamma_{i,k} r_{m,s-k} + \varepsilon_{i,s}, \quad (1)$$

where $r_{i,s}$ is the midquote return on stock i during time interval s, and $r_{m,s}$ is the return on TAIEX, Taiwan's market index. For consistency, we use the same frequencies for s as we did when computing the autocorrelation metric. We then define the R^2 from regression (1) as unconstrained, R_u^2. Next, we estimate regression (1) without the lagged market returns, effectively constraining γ to zero, and define the corresponding R^2 as constrained, R_c^2. Finally, for each stock-day, we compute:

$$price\ delay_i = 1 - \frac{R^2_{ci}}{R^2_{ui}}, \tag{2}$$

which takes values between zero and 1. A smaller delay suggests greater efficiency. Panel B of Table 3 reports the price efficiency summary statistics. To save space, here and in subsequent analyses, we report both metrics in two ways: (i) computed at the 60-second frequency and (ii) aggregated into the first principal component (PC1) across all above-mentioned frequencies. In the Internet Appendix, we report the results for additional estimation frequencies.

2.5 The control sample

The latter part of our 2019-2020 sample period is affected by the onset of the COVID-19 pandemic. We note that the pandemic influenced virtually all financial markets around the globe, but its effects are perhaps particularly comparable in regions that are geographically close to each other. In many such regions, as coronavirus cases rose, market volatility increased, and liquidity declined. As such, the true effects of the introduction of continuous trading in Taiwan may be observable if we juxtapose the TWSE against a carefully chosen control market in a DID setting. We note that since continuous trading was introduced in all TWSE stocks at the same time, the DID approach would have been prudent even in the absence of the pandemic.

As a control market, we use the Korean Stock Exchange (KRX), which is well-suited for this purpose due to its geographic proximity to the TWSE as well as its similar size. Both Taiwan and Korea faced an onset of COVID-19 cases early in the pandemic and followed similar public health strategies to contain the spread of the virus in the early 2020. These similarities allow us to cautiously claim that region-specific differences in the pandemic onset and response should not confound the DID results.

We match the TWSE and KRX firms using their market capitalization, trading volume, and volatility. Matching data come from the month prior to the sample period (that is, from October 2019), and the matching variables are estimated as described in Table 1. For comparability of market capitalization estimates, we convert share prices of TWSE and KRX stocks to the same currency. We then compute the matching score of each TWSE sample stock i and each KRX stock j as:

$$matching\ score_{ij} = \left|\frac{mcap_i}{mcap_j} - 1\right| + \left|\frac{vol_i}{vol_j} - 1\right| + \left|\frac{volat_i}{volat_j} - 1\right|, \tag{3}$$

where *mcap* is market capitalization, *vol* is the average daily volume, and *volat* is the average daily volatility. We then match, without replacement, each TWSE sample stock with the KRX stock that minimizes the matching score. When discussing the results in the following sections, we report (i) the simple TWSE-only changes in market quality variables and (ii) the DID results. The former give us a basic understanding of the changes that follow the switch to continuous trading, and the latter let us zero in on the effects attributable to the switch itself, controlling for possible confounding effects.

2.6 Regression setup

To formally examine market quality changes after the switch to continuous trading, we use DID regressions that control for trading volume, stock price volatility, and trade size. In a later section, we show that volatility and volume on the TWSE increase after the switch, while trade size declines. Hendershott, Jones, and Menkveld (2011), O'Hara and Ye (2011), Brogaard, Hagströmer, Nordén, and Riordan (2015), and Hagströmer (2021) report strong associations between these three variables and trading costs and therefore using them as controls appears warranted. For each stock i on each day t, we estimate a regression model of the following form:

$$
\begin{aligned}
DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} \\
+ \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},
\end{aligned}
\tag{4}
$$

where $DepVar_{it}$ is one of the variables of interest (e.g., the quoted spread, quoted depth, effective spread, price impact, realized spread) for stock i on day t, *Post* is an indicator variable that equals to 1 in the post-event period and zero otherwise, $TWSE$ is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks, $Volume$ is daily trading volume, $Volatility$ is the difference between the highest and lowest daily midquotes scaled by the highest midquote, and $Trade\ size$ is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample mean and divide this difference by the corresponding standard deviation. Such normalization effectively controls for fixed effects. The standard errors are double-clustered across stocks and days.

In addition to the DID, in subsequent analyses we use pre- and post-event windows that are sufficiently removed from the month of March to further reduce possible effects of pandemic-induced global volatility. To this end, our main sample period consists of a three-month pre-event window (November 2019 through January 2020) and a three-month post-event window (May through July 2020). Our results are however generally robust, as we show shortly, to including the months of February, March, and April.

3. Empirical findings

3.1 Trading cost components

Budish, Cramton, and Shim (2015) show theoretically that continuous trading decreases the ability of liquidity providers to adjust quotes in response to toxic order flow. As a result, adverse selection increases. The switch to continuous trading by the TWSE gives us a unique opportunity to test this prediction. We begin in a univariate setting by computing pre-event and post-event averages for price impacts, which serve as proxies for adverse selection of liquidity provider quotes. The results in Panel A of Table 4 suggest that the adverse selection cost increases from 4.66 bps prior to the switch to 8.43 bps post-switch.

[Table 4]

These results are consistent with the above-mentioned theoretical predictions, but they should be examined more formally to account for possible changes in liquidity cost determinants and to control for possible confounding effects. To do so, we estimate equation (4) using the price impact as the dependent variable and report the results in Panel B of Table 4. The data support the univariate findings in that adverse selection increases upon the switch to continuous trading. The interaction coefficient $Post \times TWSE$ indicates that price impacts on the TWSE increase by 0.835 standard deviations compared to the KRX. This figure translates into a 36.4% increase over the adverse selection levels that are in place during the discrete regime.

To compute the 36.4% figure, we use the standard deviations from the sample period, for which the coefficients are derived. That is, from the November 2019 through January 2020 pre-event window and the May through July 2020 post-event window. These standard deviations are computed across days rather than across stocks to better align with the event-study regression estimates. As such, the above-mentioned 36.4% value for the adverse selection increase derives from $0.835 \times 2.03 \div 4.66$, where 0.835 is the $Post \times TWSE$ coefficient estimate, 2.03 is the standard deviation, and 4.66 is the pre-event average price impact. We note that the economic magnitude of the adverse selection reduction documented here echoes that reported by Aquilina, Budish, and O'Neill (2021), who find that latency arbitrage is responsible for 33% of the total adverse selection cost on the London Stock Exchange.

To reduce the effect of volatility associated with the onset of the pandemic, our main event window contains three pre-event months (November 2019 through January 2020) and three post-event months (May through July 2020) that are removed from the month of March when it became clear that the virus had spread around the globe, multiple countries announced lockdowns, and equity prices precipitously declined. To confirm that the results are not driven by the event window choice, we repeat the analyses for two additional periods: (i) the November 2019 through July 2020 period that excludes the month of March and (ii) the entire November 2019 through July 2020 period. The results in Panel C of Table 4 are consistent with those discussed earlier and with predictions of Budish, Cramton, and Shim (2015). No matter which sample period we examine, adverse selection for the TWSE stocks substantially increases compared to their KRX matches and compared to the discreet trading regime.

While Budish, Cramton, and Shim (2015) focus on adverse selection costs, several theoretical models (e.g., Ho and Stoll (1981), Hendershott and Menkveld (2014), Aït-Sahalia and Sağlam (2017)) focus on another market maker concern – inventory management. These models recognize that holding inventory comes at a cost, and that market makers put substantial effort into mean-reverting their inventory positions. To our knowledge, the existing models do not examine the differences between inventory management in the discrete and continuous trading regimes. As such, the TWSE event may provide some background for future theoretical work.

We suggest that continuous trading may have a two-fold effect on inventory costs. On the one hand, it may allow market makers to manage their inventories more efficiently, i.e., continuously rather than once every five seconds, and as such reduce the amount of time inventory is held and the associated costs. On the other hand, and similarly to the argument made by Brogaard, Hagströmer, Nordén, and Riordan (2015), the proliferation of latency arbitrage may increase the probability of an unexpected inventory accumulation, increasing inventory costs.

Empirical studies often proxy for inventory costs using the realized spread metric. It is important to note that the metric also captures the fixed costs of technology used by market makers as well as their profits. We therefore exercise caution when interpreting the realized spread results. It is possible, for instance, that the fixed costs of market making increase in the continuous regime. Laughlin, Aguirre, and Grundfest (2014) report that market participants in the U.S. spend millions each year to maintain equipment that facilitates latency arbitrage races. In the same spirit, after the TWSE switches to continuous trading, market makers may have to invest in more advanced technology to reduce the probability of being adversely selected. Alternatively however, the increase in trading volume that (as we show shortly) is associated with continuous trading may reduce the per-share fixed cost of market making and thereby reduce realized spreads.

Finally, the switch to continuous trading may have a dual effect on market maker profits. On the one hand, the flexibility of inventory management in the continuous setting may attract new market makers. On the other hand, greater adverse selection costs and technological requirements may force some legacy market makers to exit. In the former case, enhanced competition may reduce market making profits, while in the latter case, it may increase them. The realized spreads will be affected accordingly.

Taken together, inventory costs, fixed costs, and market maker profits combine into an intricate web of effects that are difficult to disentangle using the data available to us. Nevertheless, we are able to shed light on the net impact of these effects on trading costs. On this front, we find that the realized spreads decrease. In the univariate setting, the decrease is from 17.73 to 15.07 bps, while in the multivariate setting, the decrease is by 0.505 standard deviations or 10.9%. This result suggests that continuous trading may facilitate reductions in market maker inventory costs, fixed costs, profits, or combinations of the three. As such, the continuous limit order book design appears to have both negative and positive liquidity effects, with the former manifested by the greater price impacts and the latter by the lower realized spreads. We further discuss the dual nature of these effects and the external validity of our findings in the next section.

Before moving on, we take a brief detour to discuss the transaction tax, a distinctive feature of securities trading in Taiwan. The tax is levied on all security sales, but not on purchases, at a rate of 30 bps of the sale value. Importantly, market making qualifies for a reduced rate of 15 bps. For instance, if a market maker earns 25 bps in round-trip transaction revenue, her after-tax revenue is 10 bps. The tax rate is constant throughout the sample period and therefore does not affect our main conclusions. This said, for illustrative purposes we briefly discuss market maker after-tax revenue amounts.

Panel A of Table 4 shows that the realized spreads are 17.73 bps before the switch to continuous trading and 15.07 bps after the switch. As such, the after-tax market maker revenues are 2.73 and 0.07 bps. We note that these amounts are generally consistent with figures observed in other markets. For instance, using the same 15-second estimation horizon, Conrad and Wahal (2020) show that market makers in the U.S. earn realized spreads only slightly above 0.50 bps. In the meantime, Aquilina, Budish, and O'Neill (2021) report that the realized spreads in the U.K. are on average negative, at -0.36 bps.

Are these amounts sufficient for market making to be profitable? Although this question is largely beyond the scope of our study, we point to two important considerations. First, the estimation horizon plays an important role in determining the magnitude of the realized spreads. In the main sample, we follow O'Hara (2015) and use the 15-second horizon. In the continuous regime, this horizon may be too long given that modern market makers can complete a round-trip transaction in a matter of microseconds. The Internet Appendix sheds light on the effect of shortening the horizon. For instance, when we use a 10-second horizon the post-switch after-tax realized spread increases to 1.26 bps. Further reducing the horizon to 5 seconds increases this figure to 2.91 bps. We acknowledge that such horizon reductions may not be reasonable for the entire sample period. It is unlikely that market maker inventories mean-revert within five seconds during the auction regime. Yet it is possible that the speed of mean-reversion increases when trading becomes continuous, allowing for shorter holding periods and greater realized spreads.

Second, the realized spread metric may underestimate true market making revenues. Recent work by Yao and Ye (2018) and Li, Wang, and Ye (2021) examines order submissions strategies of two investor categories: professional intermediaries and non-intermediaries. The former profit from making markets, whereas the latter build long-term positions and aim to minimize trading costs. The non-intermediaries often post aggressive limit orders undercutting the intermediaries. Doing so reduces their trading costs, since the alternative is to submit marketable orders and pay the spread. Such aggressive limit orders are likely unprofitable if submitted by the intermediaries, but are optimal for the non-intermediaries. Since our data likely contain the non-intermediary orders, the realized spreads may underestimate true intermediation revenues.

3.2 Displayed liquidity and trading costs

In a competitive market for liquidity provision, changes in trading cost components are often relayed to liquidity consumers. With this in mind, we ask if the increase in adverse selection and the decline in realized spreads post-switch affect the cost of liquidity. To shed light on this question, we examine three metrics – quoted spreads, quoted depths, and effective spreads. Quoted spreads capture displayed liquidity, that is, the difference between the lowest sell prices and the highest purchase prices posted by liquidity providers. Quoted depths quantify the number of shares available at these prices. Finally, the effective spreads account for possibilities that liquidity demanders: (i) may choose to trade when liquidity is cheaper, (ii) may occasionally receive price improvement over posted prices, and (iii) may trade at quantities that exceed those available at the best prices. In the first two cases, the quoted spreads may overstate true liquidity costs, while in the last case they may understate them.

The univariate results in Panel A of Table 5 suggest that quoted spreads and effective spreads increase, while quoted depths decline after the switch to the continuous regime. In Panel B, we examine these results in a DID regression setting with controls. Compared to the pre-event period and to the KRX stocks, quoted spreads increase by 0.508 standard deviations, or 13.3%. In turn, effective spreads increase by 0.304 standard deviations, or 5.1%. Finally, quoted depth decline by 0.193 standard deviations, or 3.4%, although this decline is only statistically significant at the 10% level. As previously, Panel C expands the event windows to first include February and April and then to include March. Both expanded window specifications suggest that quoted and effective spreads increase, and quoted depths decline after the switch.

[Table 5]

It may be of interest to discuss the economic magnitudes of these results. On the one hand, a 5.1% increase in effective spreads may appear relatively small. On the other hand, changes of similar magnitudes are rather common in the context of latency arbitrage and speed-related technologies. For instance, Shkilko and Sokolov (2020) report that when precipitation disrupts some latency arbitrage strategies in the U.S., effective spreads decline by 2.6%. Brogaard, Hagströmer, Nordén, and Riordan (2015) find that when market makers in Sweden are allowed to improve inventory management via advanced co-location technology, effective spreads decline by 2.0%. Conrad, Wahal, and Xiang (2015) show that effective spreads decline by 10% following a latency-reducing upgrade in Japan. As such, it appears that the results obtained for the TWSE align with the magnitudes reported for other markets.

Recall that the TWSE switch to continuous trading has a dual effect on trading cost components, with price impacts increasing and realized spreads decreasing. In light of these bi-directional findings, it is important to discuss the external validity of our liquidity cost results. Will the liquidity effects observed on the TWSE generalize to other markets? What are the possible consequences of mandating a switch to frequent batch auctions in developed markets such as those in Europe or the U.S.? We believe that the answer to these questions is two-fold and rather nuanced. First, the adverse selection results will likely persist. All auctions aside from the ultra-frequent should allow market makers to reprice their quotes in response to public information, thereby reducing adverse selection costs.

Second, when it comes to the realized spreads, the outcome of the above-mentioned mandate is far less certain, with the net effect depending on the balance of changes in market maker inventory costs, fixed costs, and profits. As such, the net liquidity effect of the mandate will depend on both the direction and the magnitude of realized spread changes. This discussion suggests that any market design mandates should be exercised with plenty of caution. Similarly, exchange operators who consider adopting frequent batch auctions should carefully calibrate them for the specifics of their platforms.

3.3 Price efficiency

Modern trading strategies that rely on speed and may benefit from continuous trading often improve price efficiency (e.g., Brogaard, Hendershott, and Riordan (2014), Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), Boehmer, Li, and Saar (2018)). While some of these strategies provide liquidity, others – often referred to as *toxic arbitrage* – demand it (Foucault, Kozhan, and Tham (2017)). In the discrete regime, the liquidity-taking strategies may lack profitability, as market maker quotes are not stale often enough. With the switch to continuous trading, the profitability of these strategies is likely to increase, and they may proliferate. Our earlier results are consistent with this possibility, as greater adverse selection is one possible consequence of such a proliferation. In this light, it is of interest to consider the effect of continuous trading on price efficiency. On the one hand, during the discrete regime liquidity providers may have already maintained efficiency at the optimal level by promptly adjusting their quotes. On the other hand, allowing for greater profitability of liquidity demanding strategies may have given price efficiency a boost. We examine these possibilities by turning to the price efficiency metrics.

The DID regression results in Table 6 contain only weak evidence of a relation between continuous trading and price efficiency. Out of four metrics that we examine, only one points to a price efficiency improvement, while the remaining metrics do not show significant changes. The mixed nature of the results persists when we vary estimation horizons in the Internet Appendix, with some metrics improving and others deteriorating. It therefore appears that the price efficiency benefits do not outweigh the negative liquidity effects in the wake of the TWSE switch to continuous trading.

[Table 6]

3.4 Price volatility and trading activity

In addition to greater liquidity costs, prior literature associates latency arbitrage with changes in price volatility and trading volume. Modeling a market in which liquidity takers generate toxic volume, Roşu (2019) shows that such volume is associated with greater adverse selection and greater volatility. Consistent with these predictions, Shkilko and Sokolov (2020) show empirically that latency arbitrage indeed generates substantial volume, while increasing adverse selection and volatility. Aquilina, Budish, and O'Neill (2021) report that latency arbitrage is responsible for 22% of FTSE 100 volume.

In an earlier section, we find that adverse selection indeed increases upon the switch to continuous trading. Given the above-mentioned literature, it is possible that volatility and volume increase as well. Table 7 corroborates this possibility. In both the univariate and regression settings, volatility increases. In the DID regression, volatility on the TWSE increases by 0.149 standard deviations compared to the KRX, an increase of 5.4%, corroborating the link between continuous trading and the proliferation of latency arbitrage.

[Table 7]

We next turn to trading volume, for which the effect of the switch to continuous trading may be two-fold. On the one hand, Shkilko and Sokolov (2020) show that when liquidity becomes costly, market participation by some traders (the *end users* of liquidity) declines. The end users are those who come to the market to exchange assets rather than to intermediate trade or pursue arbitrage strategies. If asset exchange becomes costly, as it does post-switch, it is possible that some end users will leave the market, and trading volume will decline. On the other hand, Aquilina, Budish, and O'Neill (2021) show that latency arbitrage generates substantial trading volume. As such, if the increase in arbitrage activity after the switch to continuous trading is sizeable, arbitrage volume may compensate for the decline in end-user volume and even result in a net volume increase.

The DID results in Table 7 suggest that the switch to the continuous regime indeed leads to trading volume increasing by 0.174 standard deviations, or 5.8%. As such, even if some end users leave the market due to the increase in trading costs, additional volume resulting from latency arbitrage outweighs the loss of end-user volume. This result is important in light of the argument in Budish, Lee, and Shim (2021), who suggest that private incentives of market operators may not be well aligned with interests of the end users of liquidity. More specifically, given that exchanges derive substantial revenue from trading fees, an increase in trading volume associated with latency arbitrage is beneficial for their profitability, even though such arbitrage may lead to greater trading costs for some market participants.

In the final step of this analysis, we turn to trade sizes, which we expect to decline for three reasons. First, the continuous regime may produce smaller trades by design, as it disaggregates liquidity demand allowing for executions at any time rather than once every five seconds. Second, our earlier results point to a decrease in quoted depth in the continuous regime, which may make large transactions less feasible. Finally, trades generated by latency arbitrage are often relatively small, as they tend to consume the leftover stale quotes. Since such trades proliferate after the switch to continuous trading, the average trade size may decrease. The results are consistent with our expectations. In the DID setting, trade sizes indeed decline.

3.5 The cross-section

In Budish, Cramton, and Shim (2015), latency arbitrage opportunities arise from publicly observable sources. These sources are likely numerous, and we rely on the existing literature for guidance on what they might be.

Goettler, Parlour, and Rajan (2009), Roşu (2019), Bhattacharya and Saar (2020), and Riccó, Rindi, and Seppi (2020) predict theoretically that information should flow into prices through a set of observable conduits such as the state of the limit order book, price history, the state of liquidity, etc. Brogaard, Hendershott, and Riordan (2014) and Kwan, Philip, and Shkilko (2021) empirically confirm these predictions and show that sophisticated market participants often trade on information obtained by monitoring the conduits. Such trading is likely to result in adverse selection. For illustrative purposes, imagine that the limit order book imbalance in stock i substantially changes at time t. Since this information is simultaneously available to many market participants, it will likely trigger latency races, with market makers trying to reprice their outstanding quotes, and snipers attempting to pick off these quotes. In a continuous market, the market maker chances of winning these races are relatively low.

We suggest that latency races should be more common in stocks that trade relatively frequently. In such stocks, updates in the observable conduits are likely to occur more often. Additionally, frequently traded assets are more likely to be involved in arbitrage between asset classes such as equities and derivatives (Shkilko and Sokolov (2020)). With this in mind, it may be of interest to examine our results in the cross-section. To do so, we divide the sample into tercile groups by trading volume. Group 1 contains the higher-volume stocks, group 2 – the medium-volume stocks, and group 3 – the lower-volume stocks.

We then re-estimate the main regression models for each group. Consistent with our expectations, the data show that while adverse selection increases in all groups, the increase is more pronounced for group 1. In Table 8, price impacts increase by respectively 0.953, 0.835, and 0.725 standard deviations, or 39%, 37%, and 33%, in higher-, medium-, and lower-volume stocks. These results echo Aquilina, Budish, and O'Neill (2021), who report that the number of latency races decreases in trading frequency.

[Table 8]

While price impacts increase across the board, the realized spreads decrease, by 0.430, 0.505, and 0.658 standard deviations. We note that these estimates are not directly comparable, because the means and standard deviations of the realized spreads differ across groups. To understand the relative magnitude of the coefficients, we convert them into percentage changes as in the earlier tests. Upon conversion, the realized spreads decline by 15%, 11%, and 8.5%. As such, the switch to continuous trading reduces the realized spreads in group 1 more than in group 3.

In an earlier section, we discuss three possible channels for the realized spread changes; the inventory channel, the fixed cost channel, and the competition channel. All three channels may help explain the differences between the groups. For instance, market makers in group 1 likely derive the most benefit from the ability to manage inventories continuously, as trades in these stocks are the most frequent, occurring every 3.8 seconds after the switch. These market makers are also more likely to benefit from a per-share reduction in fixed costs, since the trading volume in group 1 increases the most (specification 6). Finally, the reductions in the two above-mentioned costs may attract additional liquidity providers enhancing the competition channel. Unfortunately, our data are not sufficiently granular to reconcile these possibilities, so we must leave further examination for future research.

We next turn to quoted spreads, depths, and effective spreads. The increases in quoted spreads are more pronounced in group 1 than group 2 and are insignificant in group 3. The effective spreads follow a similar pattern, with somewhat lower statistical significance. Likewise, the depths decrease mainly in group 1 stocks, with a weaker effect for group 2 and no effect for group 3. As such, even though adverse selection increases across the board, the magnitude of the increase in the lower-volume stocks is sufficiently small to be offset by the decreasing realized spreads. The resulting net effect on trading costs is zero. This finding reinforces the earlier discussion of trade-offs between adverse selection and other market maker considerations. It again emphasizes the fact that optimizing market design should be viewed as a careful balancing act.

It is interesting to think of these results from the angle considered by Menkveld and Zoican (2017), who model the effects of speed acquisition by exchanges in a setting that includes fast market makers, fast snipers, and slow liquidity traders. In such a setting, the effects of speeding up the market are two-fold. On the one hand, the market makers become faster when updating quotes in response to news. On the other hand, they encounter snipers more often. The net liquidity outcome therefore depends on the balance of the two effects. If the ratio of news arrivals to liquidity trader arrivals (N/LT) is small, the second effect dominates, and liquidity deteriorates when an exchange speeds up. If the ratio is large, liquidity improves.

The switch to continuous trading by the TWSE may be viewed as a special case of speed acquisition by an exchange. It is possible that while both effects are rather weak in the discrete regime, they considerably strengthen in the continuous regime. Viewing our results through the prism of the model more literally, we suggest that N/LT on the TWSE is relatively small, and as such speeding up the market leads to a net-negative liquidity effect. In the cross-section, N/LT is smaller in group 1 compared to group 3. Per our earlier discussion, it is unlikely that the number of public news arrivals in the higher-volume stocks is smaller than that in the lower-volume stocks. It is therefore more likely that the number of liquidity traders in group 1 exceeds that in group 3.

Finally, consistent with the notion that arbitrage opportunities are more frequent in group 1, the increase in group 1 volume is greater than the increase in group 2. Meanwhile, volume does not increase in group 3. This cross-sectional result is consistent with our earlier arguments. If the number of arbitrage opportunities in group 3 is relatively small, the increase in arbitrage volume may not exceed the loss of end-user volume, resulting in a zero net effect on total volume.

Overall, the cross-sectional findings are consistent with our expectations. In frequently-traded stocks, where arbitrage opportunities are likely more abundant, the switch to continuous trading results in a greater increase in adverse selection and an increase in trading costs. In stocks that trade less frequently, the interaction of greater adverse selection and lower realized spreads delivers a net-zero effect on trading costs.

4. Conclusion

Market structure theory suggests that the continuous limit order book – the market design that dominates modern equity trading – is prone to generating adverse selection through the latency arbitrage channel. For every market maker order that may be attempting to change a stale quote based on new public information, there likely to be multiple liquidity demanding orders aiming to pick off this quote. Because the continuous limit order book processes orders one by one, the odds of replacing a stale quote before it is picked off are rather low. As such, the adverse selection cost born by market makers is high. To compensate for this cost, market makers keep the spreads wider than they would under an alternative design. Frequent batch auctions, in which orders from all market participants accumulate for a brief period of time before being matched, are often discussed as a superior alternative to the status quo.

The empirical literature has not yet examined this issue directly because transitions between market designs are rare. We examine one such recent transition, whereby a large equity market, the Taiwan Stock Exchange (TWSE), moves all of its equity trading from batch auctions to a continuous book. In a difference-in-differences setting, we find support for the above-mentioned theory predictions. Adverse selection increases significantly. Partly offsetting this increase, other trading cost components (as captured by the realized spreads) decrease. In the full sample, the adverse selection effect dominates, and trading costs increase.

Public information that triggers latency arbitrage usually comes from several sources, among which are stocks-specific channels such as changes in the limit order book and order intensity, as well as outside channels such as price changes in correlated assets and public announcements. In frequently-traded stocks, many of these channels trigger arbitrage races more often, and we expect that the adverse selection effects will be more pronounced in such stocks. The data corroborate this expectation; adverse selection indeed increases the most in stocks that trade frequently. In the less frequently-traded stocks, adverse selection also increases; however, the decline in realized spreads entirely offsets this increase, and trading costs do not change. These results further highlight the fact that optimizing market design is an intricate balancing act, with net liquidity effects dependent on multiple factors.

The increase in trading costs may affect market participation by some end-users of liquidity, that is, investors who come to the market for the purpose of traditional asset exchange rather than to engage in latency arbitrage. As trading costs increase, some end-users may choose to stay on the sidelines, and trading volume may decline. Notably, this decline may be compensated for, and even surpassed, by an increase in trading volume generated by latency arbitrageurs. The data suggest that for the frequently-traded stocks the latter effect dominates; their total trading volume increases after the switch to continuous trading.

This result offers additional empirical evidence for the ongoing debate about the costs and benefits of market design alternatives. On the one hand, the adverse selection cost embedded in the continuous design may be reduced by switching to frequent batch auctions, thus benefiting the end-users of liquidity. On the other hand, the continuous design comes with greater trading volumes boosted by arbitrage activity, thus benefiting the exchanges. Given the high fixed costs of running an exchange, it is unlikely that market operators will willingly seek to change the currently dominant continuous design.

References

Aït-Sahalia, Y., and M. Sağlam, 2017, "High frequency market making: Optimal quoting," Working paper, Princeton University.

Aquilina, M., E. B. Budish, and P. O'Neill, 2021, "Quantifying the high-frequency trading "arms race": A simple new methodology and estimates," *Quarterly Journal of Economics,* forthcoming.

Bhattacharya, A., and G. Saar, 2020, "Limit Order Markets under Asymmetric Information," *Working Paper*, University of Chicago and Cornell University.

Boehmer, E., D. Li, and G. Saar, 2018, "The competitive landscape of high-frequency trading firms," *Review of Financial Studies*, 31(6), 2227–2276.

Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan, 2015, "Trading fast and slow: Colocation and liquidity," *Review of Financial Studies*, 28(12), 3407–3443.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, "High-frequency trading and price discovery," *Review of Financial Studies*, 27(8), 2267–2306.

———, 2019, "Price Discovery without Trading: Evidence from Limit Orders," *Journal of Finance*, 74, 1621–1658.

Budish, E., P. Cramton, and J. Shim, 2014, "Implementation details for frequent batch auctions: Slowing down markets to the blink of an eye," *American Economic Review*, 104(5), 418–24.

———, 2015, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130(4), 1547–1621.

Budish, E., R. S. Lee, and J. J. Shim, 2021, "A Theory of Stock Exchange Competition and

Innovation: Will the Market Fix the Market?," Working paper, National Bureau of Economic Research.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2014, "Rise of the machines: Algorithmic trading in the foreign exchange market," *Journal of Finance*, 69(5), 2045–2084.

Chakrabarty, B., R. Pascual, and A. Shkilko, 2015, "Evaluating trade classification algorithms: Bulk volume classification versus the tick rule and the Lee-Ready algorithm," *Journal of Financial Markets*, 25, 52–79.

Conrad, J., and S. Wahal, 2020, "The term structure of liquidity provision," *Journal of Financial Economics*, 136(1), 239–259.

Conrad, J., S. Wahal, and J. Xiang, 2015, "High-frequency quoting, trading, and the efficiency of prices," *Journal of Financial Economics*, 116(2), 271–291.

Foucault, T., R. Kozhan, and W. W. Tham, 2017, "Toxic arbitrage," *Review of Financial Studies*, 30(4), 1053–1094.

Goettler, R. L., C. A. Parlour, and U. Rajan, 2009, "Informed traders and limit order markets," *Journal of Financial Economics*, 93(1), 67–87.

Haas, M., M. Khapko, and M. Zoican, 2020, "Speed and learning in high-frequency auctions," *Journal of Financial Markets*, forthcoming.

Hagströmer, B., 2021, "Bias in the effective bid-ask spread," *Journal of Financial Economics*, forthcoming.

Hendershott, T., and C. M. Jones, 2005, "Island goes dark: Transparency, fragmentation, and regulation," *Review of Financial Studies*, 18(3), 743–793.

Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, "Does Algorithmic Trading Improve Liquidity?," *Journal of Finance*, 66(1), 1–33.

Hendershott, T., and A. J. Menkveld, 2014, "Price pressures," *Journal of Financial Economics*, 114(3), 405–423.

Hendershott, T., and P. C. Moulton, 2011, "Automation, speed, and stock market quality: The NYSE's hybrid," *Journal of Financial Markets*, 14(4), 568–604.

Ho, T., and H. R. Stoll, 1981, "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial Economics*, 9(1), 47–73.

Hou, K., and T. J. Moskowitz, 2005, "Market frictions, price delay, and the cross-section of expected returns," *Review of Financial Studies*, 18(3), 981–1020.

Johann, T., T. J. Putniņš, S. Sagade, and C. Westheide, 2019, "Quasi-dark trading: The effects of banning dark pools in a world of many alternatives," Working Paper, University of Mannheim.

Kwan, A., R. Philip, and A. Shkilko, 2021, "The Conduits of Price Discovery: A Machine Learning Approach," *Available at SSRN 3710491*.

Laughlin, G., A. Aguirre, and J. Grundfest, 2014, "Information transmission between financial markets in Chicago and New York," *Financial Review*, 49(2), 283–312.

Lee, C., and M. Ready, 1991, "Inferring trade direction from intraday data," *Journal of Finance*, 46(2), 733–746.

Li, S., X. Wang, and M. Ye, 2021, "Who provides liquidity, and when?," *Journal of Financial Economics*, 141, 968–980.

Menkveld, A. J., and M. A. Zoican, 2017, "Need for speed? Exchange latency and liquidity," *Review of Financial Studies*, 30(4), 1188–1228.

O'Hara, M., 2015, "High frequency market microstructure," *Journal of Financial Economics*, 116(2), 257–270.

O'Hara, M., and M. Ye, 2011, "Is market fragmentation harming market quality?," *Journal of Financial Economics*, 100(3), 459–474.

Riccó, R., B. Rindi, and D. J. Seppi, 2020, "Information, Liquidity, and Dynamic Limit Order Markets," *Working Paper*, Bocconi University.

Roşu, I., 2019, "Liquidity and Information in Limit Order Markets," *Journal of Financial and Quantitative Analysis*, pp. 1–48.

Roşu, I., 2019, "Fast and slow informed trading," *Journal of Financial Markets*, 43, 1–30.

Shkilko, A., and K. Sokolov, 2020, "Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs," *Journal of Finance*, 75(6), 2899–2927.

Yao, C., and M. Ye, 2018, "Why trading speed matters: A tale of queue rationing under price controls," *The Review of Financial Studies*, 31(6), 2157–2183.

Table 1: Sample Characteristics

The table reports summary statistics for 100 Taiwan Stock Exchange (TWSE) stocks used in the sample. To establish a baseline, and for comparability with the main regression setup, the statistics are computed during a period prior to the switch to continuous trading: November 2019 through January 2020. *Market cap.* is market capitalization computed as the product of the number of shares outstanding and the share price. *Price* is the daily closing price in New Taiwan dollars (NTD). *Number of trades* and *Volume* are daily averages, and *Volatility* is computed for each stock-day as the difference between the highest and lowest midquotes scaled by the highest midquote. The midquote is the average between the TWSE best bid and best offer prices.

	Mean	Median	Std. Dev.	10th	90th
Market cap., NTD million	283,415	119,197	9,586	54,103	475,595
Price, NTD	174.18	62.34	489.99	14.40	333.13
Volume, share thousand	11,550	5,039	23,504	684	25,100
Volatility	0.016	0.012	0.013	0.004	0.032
Trade size, shares	7,313	4,891	8,635	2,086	13,412

Table 2: Auction Data: A Sample

The table contains a sample of the TWSE auction data for seven consecutive auctions. For each auction, the data report the allocation price, volume, as well as the bid/ask prices and sizes arising after the allocation stage. For each pair of bid and ask quotes, we compute the midquote and assign trade direction in the spirit of the Lee and Ready (1991) algorithm. That is, we compare the auction allocation price to the quotes resulting from the previous auction and assign trades executing at the bid (ask) as seller-initiated (buyer-initiated). Trades executed at the midquote are assigned direction based on the sign of the previous trade.

Auction	Price	Volume	Bid	Bid depth	Ask	Ask depth	Midquote	Trade sign
1	297.50	2,000	297.50	373,000	298.00	2,314,000	297.75	N/A
2	297.50	3,000	297.50	370,000	298.00	2,318,000	297.75	-
3	298.00	1,000	297.50	376,000	298.00	2,124,000	297.75	+
4	297.50	8,000	297.50	371,000	298.00	2,112,000	297.75	-
5	297.50	1,000	297.50	376,000	298.00	2,086,000	297.75	-
6	297.50	356,000	297.50	21,000	298.00	2,084,000	297.75	-
7	297.50	38,000	297.00	919,000	297.50	11,000	297.25	-

Table 3: Liquidity and Price Efficiency Statistics

The table reports liquidity and price efficiency statistics for 100 Taiwan Stock Exchange (TWSE) stocks used in the sample. To establish a baseline, and for comparability with the main regression setup, the statistics are computed during a period prior to the switch to continuous trading: November 2019 through January 2020. Panel A reports statistics for liquidity costs. *Quoted spread* is the difference between the best offer and the best bid. *Quoted depth* is the average of the best bid and best ask quote sizes. *Effective spread* is twice the signed difference between the traded price and the midquote immediately preceding the trade. *Price impact* is twice the signed difference between the midquote immediately preceding the trade and the midquote 15 seconds after the trade. *Realized spread* is the difference between the effective spread and price impact. To sign trades, we use the Lee and Ready (1991) algorithm. All statistics other than the quoted depths are scaled by the contemporaneous midquotes. Quoted spreads and depths are equally-weighted, and all remaining liquidity metrics are volume-weighted. Panel B reports two price efficiency metrics: return autocorrelation and price delay. *Return autocorrelation* is defined as the absolute first order midquote return autocorrelation computed at the 60-second frequency. In addition, we report the first principal component (PC1) for several estimation frequencies: 10s, 30s, 60, and 300s. *Price delay* is computed by comparing R^2s from two regressions of stock returns on market returns (equation (1)). The first (unconstrained) regression allows for several lags of market returns, while the second (constrained) model does not allow for lagged market returns (Section 2 contains estimation details). The two R^2s are then compared to compute the price delay metric as per equation (2). We report the results estimated using the 60-second frequency, and the first principal component of price delays estimated at 10-, 30-, 60-, and 300-second frequencies.

	Mean	Median	Std. Dev.	10th	90th
Panel A: Displayed liquidity and trading costs					
Quoted spread, bps.	23.05	18.84	49.31	11.37	38.78
Quoted depth, share thousand	427.73	94.84	966.93	11.24	1,023.41
Effective spread, bps.	22.38	18.61	28.15	11.56	38.49
Price impact, bps.	4.66	3.50	4.27	0.54	10.37
Realized spread, bps.	17.73	13.63	28.25	7.47	33.67
Panel B: Price efficiency metrics					
Return autocorrelation (60s)	0.09	0.07	0.09	0.00	0.22
Return autocorrelation (PC1)	0.32	0.30	0.21	0.06	0.61
Price delay (60s)	0.85	0.93	0.19	0.60	1.00
Price delay (PC1)	0.82	0.87	0.17	0.59	0.97

Table 4: Trading Cost Components

The table examines changes in price impacts and realized spreads around the move to continuous trading. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. The sample period spans November 1, 2019 to July 30, 2020. To avoid the effects of the onset of COVID-19 pandemic, in Panels A and B, the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. Panel C examines alternative event windows. Panel A contains univariate results for the TWSE stocks. Panels B and C report the results of a difference-in-differences (DID) regression of the following form:

$$DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},$$

where *DepVar* is price impact or realized spread, *Post* is an indicator variable that equals to 1 for the post-event period and zero otherwise; $TWSE$ is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks; $Volume$ is daily trading volume in stock i on day t; $Volatility$ is the difference between the highest and lowest midquotes scaled by the highest midquote; and $Trade\ size$ is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample stock mean and divide this difference by the corresponding standard deviation. As such, the model controls for stock fixed effects. Standard errors (in parentheses) are double-clustered across stock and time dimensions. *** and ** indicate statistical significance at the 1% and 5% levels. The number of observations in Panel B is 24,284.

	Price impact		Realized spread	
	[1]		[2]	
Panel A: Univariate results				
Pre	4.66		17.73	
Post	8.43		15.07	
Panel B: Regression results				
Post	0.039		0.226	***
	(0.05)		(0.09)	
TWSE	-0.436	***	0.261	***
	(0.03)		(0.06)	
Post×TWSE	0.835	***	-0.505	***
	(0.06)		(0.12)	
Volume	-0.100	***	-0.055	**
	(0.02)		(0.03)	
Volatility	0.522	***	-0.266	***
	(0.03)		(0.04)	
Trade size	0.089	***	-0.031	
	(0.02)		(0.02)	
Intercept	-0.020		-0.115	***
	(0.03)		(0.05)	
Adj. R^2	0.363		0.164	
Panel C: Regression: alternative sample periods				
Post×TWSE: excluding March	0.809	***	-0.475	***
	(0.06)		(0.10)	
Post×TWSE: full sample	0.866	***	-0.404	***
	(0.06)		(0.09)	

Table 5: Displayed Liquidity and Trading Costs

The table examines changes in displayed liquidity and trading costs around the move to continuous trading. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. The sample period spans November 1, 2019 to July 30, 2020. To avoid the effects of the onset of COVID-19 pandemic, in Panels A and B the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. Panel C examines alternative event windows. Panel A contains univariate results for the TWSE stocks. Panels B and C report the results of a DID regression of the following form:

$$DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},$$

where *DepVar* is the quoted spread, quoted depth, or effective spread, *Post* is an indicator variable that equals to 1 for the post-event period and zero otherwise; *TWSE* is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks; *Volume* is daily trading volume in stock i on day t; *Volatility* is the difference between the highest and lowest midquotes scaled by the highest midquote; and *Trade size* is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample mean and divide this difference by the corresponding standard deviation. As such, the model controls for stock fixed effects. Standard errors (in parentheses) are double-clustered across stock and time dimensions. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels. The number of observations in Panel B is 24,284.

	Quoted spread		Quoted depth		Effective spread	
	[1]		[2]		[3]	
Panel A: Univariate results						
Pre	23.05		427.73		22.38	
Post	24.48		323.48		23.50	
Panel B: Regression results						
Post	0.365	***	-0.024		0.351	***
	(0.12)		(0.09)		(0.12)	
TWSE	-0.270	***	0.101	*	-0.163	**
	(0.08)		(0.06)		(0.08)	
Post×TWSE	0.508	***	-0.193	*	0.304	**
	(0.15)		(0.11)		(0.15)	
Volume	-0.255	***	0.501	***	-0.207	***
	(0.02)		(0.03)		(0.02)	
Volatility	0.222	***	-0.400	***	0.263	***
	(0.03)		(0.03)		(0.03)	
Trade size	0.086	***	0.010		0.078	***
	(0.03)		(0.02)		(0.02)	
Intercept	-0.186	***	0.012		-0.179	***
	(0.06)		(0.05)		(0.06)	
Adj. R^2	0.137		0.137		0.114	
Panel C: Regression: alternative sample periods						
Post×TWSE: excluding March	0.514	***	-0.220	**	0.305	**
	(0.13)		(0.10)		(0.13)	
Post×TWSE: full sample	0.463	***	-0.260	***	0.294	***
	(0.10)		(0.09)		(0.10)	

Table 6: Price Efficiency

The table examines changes in return autocorrelation and price delay metrics around the move to continuous trading. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. The sample period spans November 1, 2019 to July 30, 2020. To avoid the effects of the onset of COVID-19 pandemic, in Panels A and B, the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. Panel C examines alternative event windows. Panel A contains univarate results. Panels B and C report results from a DID regression of the following form:

$$DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},$$

where *DepVar* are the autocorrelation and delay metrics for the 60-second intervals and the first principal components (PC1) of these metrics computed for 10-, 30-, 60-, and 300-second intervals, *Post* is an indicator variable that equals to 1 for the post-event period and zero otherwise; *TWSE* is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks; *Volume* is daily trading volume in stock *i* on day *t*; *Volatility* is the difference between the highest and lowest midquotes scaled by the highest midquote; and *Trade size* is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample mean and divide this difference by the corresponding standard deviation. As such, the model controls for stock fixed effects. Standard errors (in parentheses) are double-clustered across stock and time dimensions. *** and ** indicate statistical significance at the 1% and 5% levels. The number of observations in Panel B is 24,284.

	Return autocorrelation				Price delay			
	60s		PC1		60s		PC1	
	[1]		[2]		[3]		[4]	
Panel A: Univariate results								
Pre	0.095		0.319		0.852		0.818	
Post	0.100		0.252		0.739		0.719	
Panel B: Regression results								
Post	0.113	***	0.229	***	-0.542	***	-0.622	***
	(0.04)		(0.05)		(0.07)		(0.07)	
TWSE	-0.003		0.284	***	0.010		-0.016	
	(0.03)		(0.04)		(0.04)		(0.04)	
Post×TWSE	0.004		-0.548	***	-0.007		0.044	
	(0.05)		(0.08)		(0.08)		(0.09)	
Volume	0.019		0.009		0.091	***	0.114	***
	(0.01)		(0.01)		(0.02)		(0.02)	
Volatility	-0.106	***	-0.095	***	-0.128	***	-0.145	***
	(0.01)		(0.02)		(0.03)		(0.03)	
Trade Size	0.036	***	0.057	***	-0.042	***	-0.031	**
	(0.01)		(0.01)		(0.01)		(0.02)	
Incercept	-0.058	***	-0.117	***	0.276	***	0.317	***
	(0.02)		(0.03)		(0.04)		(0.04)	
Adj. R^2	0.010		0.034		0.089		0.105	
Panel C: Regression: alternative sample periods								
Post×TWSE: excluding March	-0.002		-0.549	***	0.011		0.049	
	(0.05)		(0.08)		(0.08)		(0.08)	
Post×TWSE: full sample	-0.023		-0.583	***	-0.013		0.014	
	(0.05)		(0.08)		(0.06)		(0.07)	

Table 7: Volatility and Trading Activity

The table examines changes in volatility and trading activity around the switch to continuous trading. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. The sample period spans November 1, 2019 to July 30, 2020. To avoid the effects of the onset of COVID-19 pandemic, in Panels A and B, the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. Panel C examines alternative event windows. Panel A contains univariate results. Panels B and C report the results of a pooled DID regression of the following form:

$$DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},$$

where *DepVar* is volatility, trading volume, or trade size, *Post* is an indicator variable that equals to 1 for the post-event period and zero otherwise; $TWSE$ is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks; *Volume* is daily trading volume (in thousands of shares) in stock i on day t; and $Volatility$ is the difference between the highest and lowest midquotes scaled by the highest midquote; and $Trade\ size$ is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample mean and divide this difference by the corresponding standard deviation. As such, the model controls for stock fixed effects. Standard errors (in parentheses) are double-clustered across stock and time dimensions. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels. The number of observations in Panel B is 24,284.

	Volatility		Volume		Trade size	
	[1]		[2]		[3]	
Panel A: Univariate results						
Pre	0.013		5,193		4,279	
Post	0.017		7,434		3,360	
Panel B: Regression results						
Post	0.037		0.160	***	0.441	***
	(0.05)		(0.05)		(0.08)	
TWSE	-0.078	***	-0.093	***	0.624	***
	(0.03)		(0.03)		(0.06)	
Post×TWSE	0.149	**	0.174	***	-1.206	***
	(0.06)		(0.06)		(0.10)	
Volume	0.649	***			0.219	***
	(0.02)				(0.02)	
Volatility			0.625	***	0.133	***
			(0.02)		(0.02)	
Trade size	0.084	***	0.134	***		
	(0.01)		(0.02)			
Intercept	-0.019		-0.081	***	-0.225	***
	(0.03)		(0.03)		(0.04)	
Adj. R^2	0.482		0.502		0.184	
Panel C: Regression: alternative sample periods						
Post×TWSE: excluding March	0.088		0.240	***	-1.260	***
	(0.06)		(0.06)		(0.09)	
Post×TWSE: full sample	0.127	*	0.211	***	-1.207	***
	(0.07)		(0.06)		(0.09)	

Table 8: The Cross-Section

The table examines changes in displayed liquidity, trading costs, and trading volume around the switch to continuous trading. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. To examine the effects in the cross-section, we divide the sample into terciles for higher-, medium-, and lower-volume stocks. The sample period spans November 1, 2019 to July 30, 2020. To avoid the effects of the onset of COVID-19 pandemic, the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. We report the β_3 coefficient estimates from a pooled DID regression of the following form:

$$DepVar_{it} = \alpha + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \delta_3 Trade\ size_{it} + \varepsilon_{it},$$

where *DepVar* is price impact, quoted, effective, or realized spread, quoted depth, or trading volume, *Post* is an indicator variable that equals to 1 for the post-event period and zero otherwise; $TWSE$ is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks; *Volume* is daily trading volume (in thousands of shares) in stock i on day t; $Volatility$ is the difference between the highest and lowest midquotes scaled by the highest midquote; and $Trade\ size$ is the average trade size. All continuous variables are winsorized at 1% and normalized, that is, from each stock-day observation we subtract the sample mean and divide this difference by the corresponding standard deviation. As such, the model controls for stock fixed effects. Standard errors are double-clustered across stock and time dimensions. ***, **, and * indicate statistical significance at the 1% 5%, and 10% levels.

	Price impact		Realized spread		Quoted spread		Quoted depth		Effective spread		Volume	
	[1]		[2]		[3]		[4]		[5]		[6]	
Higher	0.953	***	-0.430	**	0.715	***	-0.392	**	0.467	*	0.371	***
Medium	0.835	***	-0.505	***	0.508	***	-0.193	*	0.304	**	0.174	***
Lower	0.725	***	-0.658	***	0.170		-0.027		0.006		0.051	

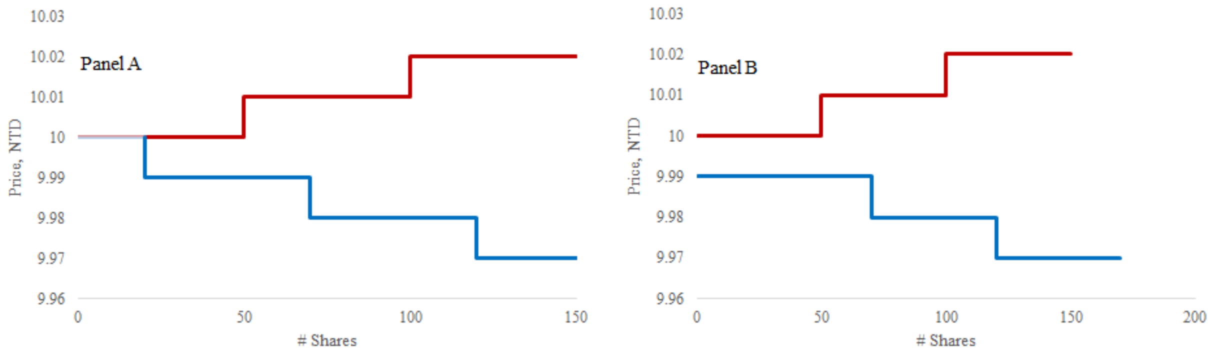

Figure 1. An auction market example

The figure plots examples of a successful and an unsuccessful auctions. In Panel A, the auctions succeeds as demand and supply cross at NTD 10.00 for 20 shares. In Panel B, the auction does not succeed, as the buyers are unwilling to pay more than NTD 9.99, while the sellers are unwilling to accept less than NTD 10.00.

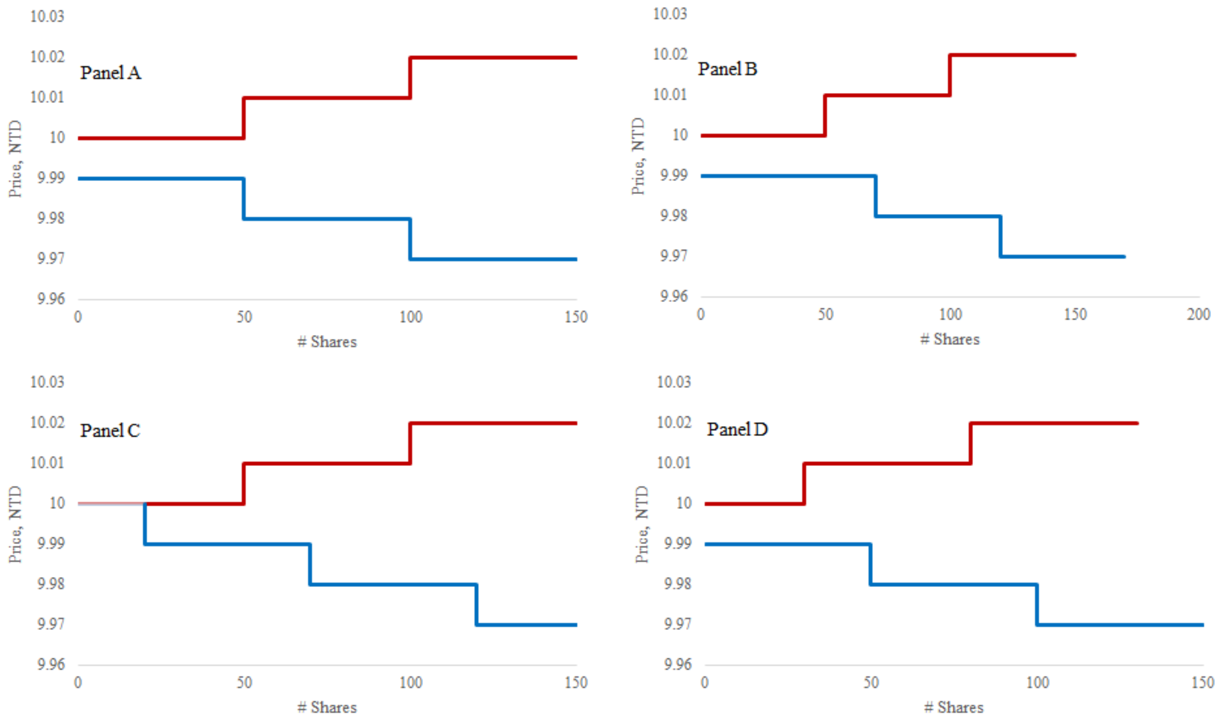

Figure 2. A continuous market example

The figure plots an example of trading in a continuous market for comparison with the auction market example in Figure 1. In Panel A, we plot supply and demand represented by resting limit orders. In Panel B, we illustrate the change in demand caused by a submission of an additional limit order to buy 20 shares at NTD 9.99. Panel C presents an alternative scenario, whereby the 20-share order to buy is marketable, and supply and demand cross resulting in a 20-share buyer-initiated trade at NTD 10.00. Finally, in Panel D we plot the state of supply and demand after the 20-share marketable order executes.