

# On the Effects of Continuous Trading<sup>\*</sup>

Ivan Indriawan<sup>†</sup> Roberto Pascual<sup>‡</sup> Andriy Shkilko<sup>§</sup>

June 3, 2021

**Abstract.** The continuous limit order book, in which messages are processed one by one in the order of receipt, is a prominent design feature of modern securities markets. Theoretical models show that this design imposes an adverse selection cost on liquidity providers and suggest that this cost may be reduced by switching to batch auctions. We examine a recent opposite move, whereby a large stock exchange switches from batch auctions to continuous trading. Consistent with theoretical predictions, we find that the move leads to greater adverse selection. The resulting increase in trading costs is associated with smaller gains from trade.

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

**JEL:** G14; G15

---

<sup>\*</sup>We thank Michael Brolley, Benjamin Connault, Pankaj Jain, Kai Li, Jane Liu, Esen Onur, John Shim, Konstantin Sokolov, Marius Zoican, and seminar/session participants at the Financial Markets and Corporate Governance Conference, Mid-Atlantic Research Conference in Finance, Midwest Finance Association annual meeting, Stern/Salomon Microstructure Conference, Macquarie University, and Wilfrid Laurier University for valuable comments. Pascual acknowledges support of the Spanish DGICYT project ECO2017-86903-P and 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>.

<sup>†</sup>Auckland University of Technology, New Zealand, e-mail: [ivan.indriawan@aut.ac.nz](mailto:ivan.indriawan@aut.ac.nz)

<sup>‡</sup>University of the Balearic Islands, Spain, e-mail: [rpascual@uib.es](mailto:rpascual@uib.es)

<sup>§</sup>Wilfrid Laurier University, Canada, e-mail: [ashkilko@wlu.ca](mailto:ashkilko@wlu.ca)

# 1. Introduction

The majority of modern equity 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 increase the level of adverse selection (toxicity), because it reduces the ability of liquidity providers to reprice stale quotes before they are picked off (Budish, Cramton, and Shim (2015)). Trading costs increase as a result. As a remedy, the literature proposes replacing continuous trading with frequent batch auctions, in which orders accumulate for a period of time before being matched against each other, thus giving market makers a better 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 decision by the Taiwan Stock Exchange (TWSE) to move all of its activity from batch auctions to continuous trading. In a difference-in-differences (DID) setup, we find that continuous trading is associated with significantly greater adverse selection, a sizeable reduction in displayed liquidity, and an increase in trading costs.

While important in their own right, changes in liquidity costs are also notable in their potential to affect investor welfare (Biais, Foucault, and Moinas (2015)). When market participation becomes costlier, some investors (the end-users of liquidity) may choose to stay on the sidelines, and gains from trade may decline. To examine this possibility in our setting, we measure trading volume generated by the uninformed liquidity seekers, a trader category that is perhaps the most sensitive to liquidity costs, and find that it declines after the switch to continuous trading. Insofar as this trader category is representative of some end-users of liquidity, the decline in its market participation is consistent with a partial reduction in gains from trade.

Notably, although volume generated by the uninformed traders declines after the switch to

continuous trading, toxic volume substantially increases, for a net increase in total volume. This result highlights an important tension in modern markets. On the one hand, exchange revenues are volume-dependent, so exchange operators have a preference for market designs that maximize volume. On the other hand, such designs may not benefit all market participants, potentially leading to reductions in gains from trade. Consistent with [Budish, Lee, and Shim \(2020\)](#), private-market incentives may therefore be insufficient to maximize end-user welfare.

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<sup>th</sup>) to such markets as the Toronto Stock Exchange (13<sup>th</sup>) and the Australian Securities Exchange (20<sup>th</sup>). 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. Notably, the data contain clean and sizeable regime shifts on the day of the switch. For instance, [Figure 1](#) shows that effective spreads, our main trading cost metric, increase sharply on March 23 and stabilize at the new level thereafter. This pattern alone may allay concerns with the confounding effects; however, in subsequent analyses we rely on a formal two-pronged approach to mitigate these concerns even further.

[Figure 1]

First, we use a DID setup with a control sample of stocks trading on the Korean Stock Exchange (KRX). The similarities in infection emergence and 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 to assess

the sensitivity of our results to possible pandemic effects. The results are preserved regardless of event window lengths and their proximity to the March 23 launch date. 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.

The 21<sup>st</sup> 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 Rordán (2015)). Market participants responded to these changes by adopting the latest technology in a speed race to the exchange engines and between markets (Baron, Brogaard, Hagströmer, and Kirilenko (2019), 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 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, she can do so even if the other traders are a little faster. Given this advantage, [Budish, Cramton, and Shim \(2015\)](#) propose that market operators should reduce their reliance on the continuous design.

While [Budish, Cramton, and Shim \(2015\)](#) focus on adverse selection costs, [Aït-Sahalia and Sağlam \(2017\)](#) examine a different market maker concern – inventory management. In their model, the market makers’ decisions are characterized by an inventory penalty function, whereby holding inventory comes at a cost. If the market maker can predict future liquidity demand more accurately, she may reduce the risk of taking on unwanted inventory and therefore the penalty cost. Empirical research corroborates this prediction. [Brogaard, Hagströmer, Nordén, and Riordan \(2015\)](#) find that a better ability to predict incoming order flow is associated with lower inventory costs, while [Shkilko and Sokolov \(2020\)](#) suggest that exposure to toxic order flow affects this predictive ability negatively. Following this line of reasoning, continuous trading may have a two-pronged effect on market making costs, by increasing both adverse selection and the risk of unexpected inventory accumulation.

Our analyses support these expectations. In the DID regression setup, we find that adverse selection on the TWSE substantially increases after the switch to continuous trading. Realized spreads too increase consistent with an increase in inventory costs. The total effect is an increase in effective spreads, our proxy for liquidity costs, and a reduction in displayed liquidity represented by quoted spreads and depths. The data also show that continuous trading brings mild improvements in price efficiency, although these results are not always statistically significant, and their economic magnitude appears secondary to that of the liquidity effects.

Finally, to shed light on the effects of continuous trading on price discovery, we examine the flow of information into prices. [Campbell, Ramadorai, and Schwartz \(2009\)](#) and [Weller \(2018\)](#) show that research into firm fundamentals facilitates gradual incorporation of earnings information prior to earnings announcements. It is possible that the higher liquidity costs that follow the

switch to continuous trading negatively affect this process. Alternatively however, the increase in liquidity costs may not be sufficiently large to affect fundamental research, as its profitability is likely greater than that of uninformed trading mentioned earlier. The results are consistent with the latter possibility; we find no changes in the pre-announcement price adjustments. As such, fundamental information appears to flow into prices at the same pace as before the switch.

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 recently obtained approval from the Securities and Exchange Commission (SEC) to implement batch auctions on one of its smaller equity exchanges, BYX.<sup>1</sup> Our results help explain the general reluctance of the industry to change the status quo. We show 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, the 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. Similar concerns may accompany analyses of recent introductions of periodic auctions in Europe. Furthermore, it should be noted that European auction mechanisms are characterized by a limited degree of transparency (e.g., Johann, Putniņš, Sagade, and Westheide (2019)) further confounding design comparisons. In the meantime, the TWSE transition to continuous trading occurs in a market characterised by a high degree of consolidation and without an accompanying change in transparency.

---

<sup>1</sup>“Cboe Receives Regulatory Approval to Launch Periodic Auctions for U.S. Equities Trading,” March 29, 2020 (<http://bit.ly/206wI61>).

An issue that deserves additional discussion is the external validity of our results. Will the liquidity effects observed on the TWSE generalize to other markets? We believe that the answer is largely *yes*, as the TWSE continuous platform closely resembles (if not fully replicates) those of the world's leading exchanges. The switch to continuous trading is the most recent, and perhaps the last, step in the multi-year process of modernizing the TWSE. In the past decade, the exchange invested heavily in platform upgrades and began offering such staples of modern trading infrastructure as high matching engine throughput, direct access to data feeds, and subscription-based colocation. All these services were in place prior to our sample period. As such, we are sufficiently confident that the TWSE is a uniquely suitable laboratory for examining the effects of moving between continuous trading and frequent batch auctions.

Before moving on, we acknowledge that our results may not generalize to a small group of ultra-high auction frequencies. [Haas, Khapko, and Zoican \(2020\)](#) argue that at such frequencies the fastest traders may, as in the continuous market, still have an edge on market makers. Given modern trading speeds, such auction frequencies likely measure in microseconds or even sub-microseconds and therefore represent a rather small group of possible designs. In the meantime, the auction frequencies that are currently under consideration or may be under consideration in the near future are several orders of magnitude greater than the ultra-high. For instance, the BYX proposal discussed earlier aims to run the auctions 10 times per second. As such, we believe that our findings are well-suited to shed light on the effects of switching to the majority of action frequencies, particularly those under consideration today.

Taken together, our results show that the continuous limit order book design is associated with greater liquidity costs, which negatively affect market participation by the uninformed investors. In the meantime, the design benefits exchanges by substantially increasing volume generated by the short-term informed traders such as latency arbitrageurs. This activity appears to moderately improve price efficiency. Overall, for these effects to be consistent with welfare maximization, society must heavily discount increases in trading costs and value increases in price

efficiency exceptionally highly. In addition, an abundance of latency arbitrage opportunities encourages market participants to engage in a rather costly speed race, devoting large amounts of capital to shaving micro- or even nanoseconds off information processing and order submission speeds (Shkilko and Sokolov (2020)). By reducing the number of latency arbitrage opportunities, frequent batch auctions may allow for capital redeployment to alternative, and potentially more socially beneficial, uses.

## 2. Related literature

Theoretical comparisons between the discreet and continuous designs trace back to Kyle (1985), who shows that a sequence of call auctions, compared to just one auction, may result in greater losses for the uninformed traders. Madhavan (1992) models informational asymmetries and shows that they are more severe in the continuous design than in the auction design. More recently, Budish, Cramton, and Shim (2015) take the design comparison to the modern marketplace characterized by continuous limit order books. They argue that continuous trading comes with substantial adverse selection and propose to revert back to frequent batch auctions. Budish, Lee, and Shim (2020) discuss competition between market operators and show that the competitors do not always have the incentives to change market design, even to the one that is more welfare-enhancing.

Switches from discreet to continuous trading have occurred previously, most recently in the 20<sup>th</sup> century. Empirical studies examine several such switches and find that they lead to outcomes different from the ones on the TWSE. Amihud, Mendelson, and Lauterbach (1997), Muscarella and Piwowar (2001), and Henke and Lauterbach (2005) examine transitions from call auctions to continuous trading on, respectively, the Tel Aviv Stock Exchange, Paris Bourse, and the Warsaw Stock Exchange. In all three cases, continuous trading results in liquidity improvements.

A direct comparison between our results and those in the above-mentioned studies is un-

fortunately impossible, since these studies use relatively coarse data and therefore must rely on indirect liquidity metrics such as the Amivest ratio or the Roll (1984) measure. This said, we believe that our results are likely complementary to these early findings, with the difference in liquidity outcomes arising because of the very infrequent nature of the auctions that they examine. Before the switch to continuous trading, the auctions in Tel Aviv occur once daily, and auctions in Paris and Warsaw occur twice daily – quite infrequently compared to the five-second auctions on the TWSE. We believe that such low auction frequencies may come with substantial inventory costs, and the reduction in these costs upon the switch to continuous trading may lead to the overall liquidity improvement.

To elaborate, recall that in our setting the switch to continuous trading has a dual effect; both adverse selection and inventory costs increase. Although the three above-mentioned studies are unable to decompose liquidity costs in the same way we do, knowing the frequencies of the auctions that they examine allows us to hypothesize a possible outcome of such a decomposition. First, we believe that the adverse selection effects in these three markets could be similar to ours. As we discuss earlier, all auctions aside from the ultra-frequent should allow market makers to reprice their quotes in response to news, thereby reducing adverse selection costs. Second, and more importantly, we posit that the inventory costs may dominate the overall cost of liquidity provision in low-frequency auctions because such auctions allow for inventory rebalancing only once or twice a day. As such, even though adverse selection costs in Tel Aviv, Paris, and Warsaw may have increased upon the switch to continuous trading, the decline in inventory costs could have eclipsed this effect leading to a net liquidity improvement.

Liquidity cost results consistent with ours have been so far documented only in experimental markets. Schnitzlein (1996) and Theissen (2000) show that trading costs are lower in an experimental auction environment compared to a continuous environment. Although informative, these experiments do not perfectly replicate actual markets. To remain tractable, experimental settings often rely on simplifying features such as restricting participants to submitting only market or-

ders and assuming that the market is solely quote-driven. Our results complement experimental evidence in that they apply recent theoretical predictions to a modern high-speed order-driven market and also, for the first time in the literature, examine the economic magnitude of the effects that accompany a transition between the two market designs.

In a recent independent study, [Riccò and Wang \(2020\)](#) also investigate the impact of the TWSE switch to continuous auctions. They too report that the switch is associated with greater spreads and trading volume, but do not examine what drives these changes. Our study differs along several dimensions. First, we examine the drivers of the aforementioned liquidity changes, that is, adverse selection and inventory costs, and tie them to recent theory. Second, we examine gains from trade and show that they decline at least for some market participants upon the switch to continuous trading. Finally, we examine price efficiency and price discovery changes that accompany the switch. As such, our study provides a comprehensive view of market quality and welfare implications of continuous trading.

Finally, this study is related to, yet distinct from, research that examines opening and closing auctions (OCAs). [Pagano and Schwartz \(2003\)](#) and [Barclay, Hendershott, and Jones \(2008\)](#) show that by consolidating supply and demand at one point in time OCAs allow for more efficient prices and more accurate price discovery. [Ellul, Shin, and Tonks \(2005\)](#) also find that OCAs benefit price discovery, but show that they suffer from high failure rates when information asymmetries are high. [Comerton-Forde and Rydge \(2006\)](#) point out the importance of OCA transparency, and [Pagano, Peng, and Schwartz \(2013\)](#) study spillover effects of the auctions on continuous trading sessions that occur during the trading day. In contrast to this stream of research, we study relatively frequent auctions that occur during the trading day and focus on liquidity as well as price efficiency and price discovery effects of replacing such auctions with continuous trading.

### **3. Data and metrics**

#### **3.1 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 from November 2019 through July 2020. To establish a baseline, Table 1 reports summary statistics computed prior to the switch to continuous trading.

The average sample stock has a market capitalization of 282 billion New Taiwan dollars (NTD), share price of NTD 182, daily volume of about 9.6 million shares, and daily volatility of 1.43 bps. We compute volatility as the difference between the highest and lowest daily midpoints scaled by the average midpoint. The sample covers a broad cross section, with market capitalizations ranging between NTD 54 billion and 474 billion (respectively, in the 10<sup>th</sup> and 90<sup>th</sup> percentiles), prices ranging between NTD 14.73 and 372.05, and daily volumes – between 0.55 and 23.1 million shares.

[Table 1]

#### **3.2 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 that is similar to that of the Trade and Quote Database often used to examine liquidity in

the U.S. The data contain all 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 omit the first and last five minutes of the 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 quote midpoint 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 quote midpoint at the time of the trade and the midpoint 30 seconds after the trade. Finally, to gauge inventory costs we follow Brogaard, Hagströmer, Nordén, and Riordan (2015) and use the *realized spread*, the difference between the effective spread and price impact.

We drop instances when the best quotes are locked or crossed, that is when the quoted spread is zero or negative. 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. All variables are scaled by the corresponding quote midpoints. In a later section, we show that the results are robust to varying horizons for price impact and realized spread estimates between 10 and 300 seconds.

### **3.3 Liquidity metrics in the discrete regime**

To describe the data available to us during the auction regime, we begin with a depiction of a 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.

Panel A of Figure 2 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 2]

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 available to sell and to buy at the best prices that remain unexecuted (respectively, 30 shares at NTD 10.00 and 50 shares at NTD 9.99). For unsuccessful auctions, the data report only the above-mentioned item (iii). For the example of an 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 an informative comparison. To explain our approach, we begin with an example from continuous markets that makes use of supply and demand curves. In Panel A of Figure 3, 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 3]

In this continuous example, the option to join the queue of bids results in a scenario similar to the unsuccessful auction in Figure 2. 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 2, this reasoning would point to a quoted spread of NTD 0.01 and a quote midpoint of NTD 9.995.

Although our approach to the auction spread estimation is somewhat novel, it is quite 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 on to the last remaining ingredient for trading cost analyses – trade signing. We suggest that methodologies that are well-established in continuous markets, specifically the Lee-Ready algorithm that benchmarks against the pre-trade quotes, 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, executing every five seconds. The midquote is stable during the first six auctions at NTD 297.75 and decreases to NTD 297.25 as the 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 auction 6, the book has 21,000 remaining shares on the bid at NTD 297.50. During auction 7, 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 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.

In the robustness section, we examine the sensitivity of our results to the trade signing algorithm by examining alternatives to Lee-Ready. Specifically, we consider the spectrum of possible trade signing approaches, from the case that assumes that most of price discovery is done through trades to the case that assumes the opposite – that price discovery is driven by quotes. No matter the approach, the main results remain intact; adverse selection and trading costs increase after the switch to continuous trading.

One remaining issue related to liquidity metrics is aggregation. In continuous markets, re-

searchers usually time-weight quoted spreads and depths when computing daily aggregates. Time-weighting is however not very informative in the discrete regime when quoted spreads are reported once per auction, that is once every five seconds. To allow for better comparability, we equal-weight quoted spreads 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, 23.41 and 19.12 bps, while price impacts and realized spreads are 10.84 and 8.27 bps. Quoted depth is about 448 thousand shares, or 4.7% of daily trading volume. Again, we observe non-trivial variation in the cross-section, with effective spreads for instance ranging from 10.16 bps in the 10<sup>th</sup> percentile to 33.89 bps in the 90<sup>th</sup> percentile, and realized spreads ranging from 0.04 to 18.74 bps.

[Table 3]

### 3.4 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 midpoint returns should have zero autocorrelation. It is defined as the absolute first order midpoint return autocorrelation, and we compute it at several frequencies  $s \in \{10s, 30s, 60s, 300s\}$ . 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  $i$ :

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

where  $r_{i,s}$  is the quote midpoint 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  $\gamma$  to zero, and define the corresponding  $R^2$  as constrained,  $R_c^2$ . Finally, for each stock-day  $i$ , we compute:

$$price\ delay_i = 1 - \frac{R_c^2}{R_u^2}, \quad (2)$$

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 robustness section, we show that our results are robust to varying horizons for both metrics.

### 3.5 The control sample

The latter part of our 2019-2020 sample period coincides with the COVID-19 pandemic. To verify that the results are not driven by the pandemic, we use the DID approach. Specifically, we surmise that the pandemic affected volatility in most equity markets in a similar way. As such, the true effect of the introduction of continuous trading in Taiwan may be observable if juxtaposed against a control market. We note that since continuous trading was introduced for all stocks simultaneously, a 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 similar size. Both Taiwan and Korea faced an onset of COVID-19 cases early in the pandemic and followed similar public health strategies managing to contain the spread of the virus in the spring of 2020. These similarities allow us to cautiously claim that region-specific differences in the pandemic onset and response should not confound the DID results.

To match the TWSE and KRX stocks, we use trading volumes and closing prices converted to the same currency for comparability. We then compute the matching score of each TWSE sample stock  $i$  and each KRX stock  $j$  as:

$$MS_{ij} = \left| \frac{P_i}{P_j} - 1 \right| + \left| \frac{V_i}{V_j} - 1 \right|, \quad (3)$$

where  $P$  is the daily average closing price, and  $V$  is the daily average dollar volume. We then match, without replacement, each TWSE sample stock with the KRX stock that minimizes the matching score. In the following sections, we report (i) the simple TWSE-only differences in market quality variables and (ii) the DID results. The former give us an understanding of the economic magnitude of 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 global confounders.

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 and December 2019 as well as January 2020) and a three-month post-event window (May-July 2020). Our results are however generally robust, as we show shortly, to including the months of February, March and April. As an illustration, Figure 4 plots trading costs from the main event window for both the TWSE and KRX.<sup>2</sup>

---

<sup>2</sup>As is conventional, we drop trading days when either the TWSE or KRX is closed for a holiday. Given the Lunar New Year celebrations that occurred in late January 2020 and several other holidays, the pre-event window is

[Figure 4]

## 4. Empirical findings

### 4.1 Adverse selection

Budish, Cramton, and Shim (2015) show theoretically that continuous trading decreases the ability of liquidity providers to adjust their 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 by computing simple pre- and post-event averages for price impacts, which serve as proxies for adverse selection of liquidity provider quotes. 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 results from alternative windows later in this section. The univariate results in Panel A of Table 4 suggest that adverse selection increases by 27%, from 10.84 bps prior to the switch to continuous trading to 13.78 bps post-switch.

[Table 4]

These results are consistent with the above-mentioned theoretical predictions, however, their univariate nature comes with caveats. First, the univariate analysis does not account for the effects of known adverse selection determinants such as trading volume and volatility. Second, they may be subject to confounding events, particularly the effects of the COVID-19 pandemic. To examine the adverse selection effects more formally, we use the following DID regression setup for each

---

four days shorter than the post-event window.

stock  $i$  on each day  $t$ :

$$\begin{aligned} price\ impact_{it} = & \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} \\ & + \delta_2 Volatility_{it} + \epsilon_{it}, \end{aligned} \quad (4)$$

where  $Post$  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, and  $Volatility$  is the difference between the highest and lowest midpoints scaled by the average midpoint. 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.

The results in Panel B of Table 4 support previously reported univariate findings in that adverse selection increases upon the switch to continuous trading. In specification 1, the DID specification without the volume and volatility controls, the interaction coefficient  $Post \times TWSE$  indicates that price impacts on the TWSE increase by 0.460 standard deviations compared to the KRX, a notable 24% increase over the adverse selection levels that are in place during the discrete regime.<sup>3</sup> In specification 2, which controls for volume and volatility, the interaction coefficient suggests that price impacts increase by 8%.<sup>4</sup>

We note that although the volatility and volume controls do not reduce statistical significance of the  $Post \times TWSE$  coefficient, they reduce its economic magnitude. On the one hand, this may suggest that some of the increase in adverse selection is attributable to changes in volume and

---

<sup>3</sup>To compute the economic significance of regression coefficients, we use standard deviations from the sample period, for which the coefficients are derived. For instance, the standard deviation for price impacts used to gauge economic significance in Panel B of Table 4 is 5.68. This estimate is from the November 2019 through January 2020 pre-event window and the May through July 2020 post-event window.

<sup>4</sup>We note that the  $Post \times TWSE$  coefficient captures the difference between the post-switch effects on the TWSE and the KRX. To measure the full economic effect for the TWSE, one should add the coefficients for  $Post$  and  $Post \times TWSE$ . Given that the  $Post$  coefficient in specification 1 is statistically indistinguishable from zero, we base the economic interpretation on the  $Post \times TWSE$  coefficient alone. In specification 2, in which the  $Post$  coefficient is significant, we use  $Post + Post \times TWSE$ .

volatility, the two known adverse selection determinants. In a subsequent section, we show that both of these determinants increase upon the switch to continuous trading. On the other hand, the price impact, volume, and volatility are all subject to the same structural break that occurs on the day of the switch. As such, the two control variables may mechanically subsume some variation in price impact. While it is not possible to gauge which of the two effects dominates, we suggest that the coefficient in specification 2 likely represents the lower bound of the economic effect, while the coefficient in specification 1 represents the upper bound. In subsequent discussions, we focus on the lower bound coefficients to remain conservative.

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 markets 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. 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.

## **4.2 Displayed liquidity and trading costs**

Adverse selection is a cost of market making. In competitive markets, changes in this cost are often relayed to liquidity consumers. With this in mind, we now ask if the increase in adverse selection post-switch affects the cost of liquidity. To answer this question, we examine two related metrics – quoted and effective spreads. The former captures displayed liquidity, that is, prices posted by liquidity providers. The latter accounts for two additional possibilities: (i) that liquidity

demanders may choose to trade when liquidity is cheaper, and (ii) that they occasionally receive price improvement over posted prices.

The univariate results in Panel A of Table 5 indicate that quoted spreads increase and quoted depths decline after the switch to the continuous regime. In Panel B, we confirm these results in a DID regression setting of equation (4). Compared to the pre-event period and to the KRX stocks, quoted spreads increase by 0.907 standard deviations, equivalent to 14%. Another notable change is the 0.380 standard deviations decline in quoted depth, equivalent to 10% of the pre-switch depth figure.

In Table 6, we expand the DID regression analysis to effective and realized spreads. Effective spreads, which capture the cost of taking liquidity, increase by 1.149 standard deviations, equivalent to 21%. Notably, this figure is close to the estimate in Aquilina, Budish, and O’Neill (2020), who study latency arbitrage races in the continuous limit order book environment and show that eliminating such races should improve liquidity by about 17%. Next, we turn to the realized spreads that are a composite metric used to proxy for market maker inventory costs. Brogaard, Hagströmer, Nordén, and Riordan (2015) and Shkilko and Sokolov (2020) show that unpredictable order flow such as that generated in the process of latency arbitrage may impede market maker inventory management. When arbitrageurs pick off stale quotes, market maker inventory may change unexpectedly, requiring additional efforts to balance it. Inventory holding costs increase as a result. The results corroborate this possibility. Panel B of Table 6 shows that realized spreads increase by 0.545 standard deviations upon the switch to continuous trading. The results for the two alternative sample periods in Panel C are consistent with these findings.

[Tables 5 and 6]

Before moving on, it is useful to discuss two issues related to realized spreads. As a residual metric (the difference between effective spreads and price impacts), realized spreads capture not only the inventory costs, but also order processing costs and liquidity provider profits. Our dis-

cussion of this metric has so far focused solely on the inventory costs. We cautiously suggest that this focus is justified given that it is difficult to think of ways, in which continuous trading would increase order processing costs per share. If anything, given the greater volumes resulting from continuous trading and the fact that order processing costs have a non-trivial fixed component, these costs could have declined.<sup>5</sup> When it comes to profits, it is again difficult to think of a scenario, in which these could appreciably change in a competitive market for liquidity provision. One possibility is that the switch to continuous trading forced some market makers to exit, resulting in a less competitive environment and therefore greater per-share profits. Nevertheless, a media search and conversations with industry participants do not produce any evidence of market maker exits as a result of the switch.

### 4.3 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

---

<sup>5</sup>We formally discuss increases in trading volume shortly.

a boost. We examine these possibilities by turning to the price efficiency metrics.

Table 7 shows that the effects of continuous trading on price efficiency are somewhat mixed. First, the autocorrelation metric and the principal component of this metric suggest that price efficiency improves, with the DID coefficients of -0.181 and -0.056, respectively. It should be noted that this improvement is economically moderate, between 1.4% and 3.3%. Second, the DID coefficients for the price delay metric are -0.215 and -0.058, translating to improvements between 0.4% and 2.0%. Notably however, changes in the price delay metric are mostly insignificant when we vary the estimation window in Panel C, making price delay the only metric so far that does not show stable results across estimation windows. As such, it appears that continuous trading moderately improves some, but not all, aspects of price efficiency.

[Table 7]

In light of these results, it may be of interest to contemplate the net effect of continuous trading. On the one hand, reductions in return autocorrelations, even on the level of 3.3%, benefit market participants by increasing the probability of trading at the most up-to-date prices. On the other hand, this benefit comes at a cost to liquidity. Consistent with Foucault and Moinas (2019), to justify this tradeoff as welfare-enhancing the benefits of relatively small improvements in price efficiency must be sizeable, and traders must value them exceptionally highly.

#### **4.4 Volatility, volume, and gains from trade**

In this section, we seek to better understand the effects of continuous trading on gains from trade. To proceed, we first outline the links between latency arbitrage, volatility, and trading volume proposed by recent theoretical and empirical work and then examine these links in our setting. Modeling a market in which liquidity takers generate toxic volume, Roşu (2019) shows that such volume is associated with increased adverse selection and volatility. Consistent with these predictions, Shkilko and Sokolov (2020) show empirically that liquidity-taking latency arbitrage

indeed generates substantial volume, while increasing adverse selection and volatility. In an earlier section, we find that adverse selection increases upon the switch to continuous trading and relate this increase to the proliferation of latency arbitrage. Given the above-mentioned literature, it is possible that volatility increases as well. We examine this possibility in Table 8. In the DID setting, volatility indeed increases by 0.175 standard deviations after the switch (specification 2). We note that, aside from its standalone significance, this result justifies our use of volatility as a control in all regression specifications.

[Table 8]

We next turn to the volume effects. The theoretical literature emphasizes the role of liquidity in promoting welfare. Improved liquidity allows greater numbers of economic agents to come to the market and benefit from exchanging assets, increasing gains from trade. When liquidity is costly, some agents (we call them the *traditional users* or *end-users* of liquidity) may choose to stay on the sidelines, and gains from trade are reduced. Since the switch to continuous trading results in greater liquidity costs, it is possible that some end-users will leave the market, and trading volume will decline. Still, if the increase in arbitrage activity is substantial, arbitrage volume may compensate for this decline and even result in a net volume increase.

We begin to examine these possibilities in Table 8. At first glance, the univariate results in Panel A and the regression results in specification 3 of Panel B suggest that the switch to continuous trading leads to a volume increase. Notably however, when we control for volatility in specification 4, the change in volume becomes insignificant. This latter result is noteworthy. Insofar as changes in volatility proxy for the proliferation of latency arbitrage discussed by Roşu (2019) and Shkilko and Sokolov (2020), the latter result is consistent with the notion that continuous trading may not lead to greater gains from trade for the traditional users of liquidity.

To examine this issue further, we focus on trades that generate the least adverse selection. We reason that if higher post-switch trading costs were to force some end-users out of the market, the

users affected the most would be those, who are the least informed. For them, market participation is likely costlier than for those who trade on superior information since trading costs of the informed traders are offset by greater trading profits. An increase in the cost of trading should therefore affect the uninformed traders the most. To proxy for the volume generated by such market participants, we take all trades that have negative price impacts and refer to trading volume resulting from these trades as *low-toxicity volume*.

Panel A of Table 9 shows that such volume declines substantially after the switch to continuous trading, from nearly 1.2 to about 0.4 million shares per day. As previously, we next examine changes in low-toxicity volume in the DID regression setting. Panel B of Table 9 confirms the univariate findings in that low-toxicity volume declines after the switch to continuous trading. To reiterate, insofar as low-toxicity volume captures market participation by some of the traditional users of liquidity, these results point to a possible reduction in gains from trade.

[Table 9]

Two caveats should be noted in light of the above-mentioned results. First, when liquidity becomes more expensive, the uninformed may switch from demanding to supplying it. If so, even though they do not initiate as many non-toxic transactions as before, their market participation may not necessarily decline. Although plausible, this narrative does not easily reconcile with our earlier results, as such additional liquidity supply should result in tighter spreads and greater quoted depths – the effects opposite to those we find in the data.

Second, recall that while low-toxicity volume declines, total volume increases (Table 8), pointing to an increase in toxic volume. So far, we have assumed that latency arbitrage comprises the lion's share of toxic volume. An alternative, however, is that toxic volume increases due to fundamentally informed investors arriving to the market more frequently or choosing to trade more aggressively. Either way, an increase in their participation may be beneficial for price discovery and as such represent a positive outcome of the switch to continuous trading. Our data are

not sufficiently detailed to study this possibility directly, as we do not have access to trader accounts, so in what follows we examine indirect evidence on market participation by the informed.

To do so, we rely on the approach used by [Weller \(2018\)](#), who suggests that greater informed investor activity should allow the market to discover more earnings news prior to earnings announcements. As market participants research firm fundamentals, value-relevant information flows into prices through their trading. In the case of earnings, the more information is discovered prior to an announcement, the smaller should be the market reaction to the announcement itself. To measure this effect, [Weller \(2018\)](#) uses the price jump ratio, PJR, that divides the earnings announcement return by the total return plausibly attributable to the announcement. The latter includes three weeks of pre-announcement price changes. A low PJR is consistent with high levels of price discovery, with a substantial portion of earnings information incorporated into prices in the weeks prior to the announcement. In our setting, if the presence of fundamentally informed traders indeed increases after the switch to continuous trading, PJR should decline. To compute PJR, we follow [Weller \(2018\)](#) and let  $T$  be the earnings announcement date. We then define the *announcement window* as  $[T - 1, T + 2]$ , *event window* as  $[T - 21, T + 2]$ , and *pre-event window* as  $[T - 255, T - 90]$ . For each day  $t$  and each stock  $i$ , we compute the close-to-close return,  $r_{it}$ , and the return on each region's market index,  $r_{mt}$ . We then obtain the abnormal return,  $abr_{it}$ , as the difference between the stock  $i$  return on day  $t$  and the expected return according to the market model estimated in the pre-event window, that is,

$$abr_{it} = r_{it} - \hat{\alpha}_i - \hat{\beta}_i r_{mt}. \quad (5)$$

Next, we define cumulative abnormal return as the sum of abnormal returns from  $t_1$  to  $t_2$ ,

$$CAR_i^{t_1, t_2} = \sum_{t=t_1}^{t_2} abr_{it}, \quad (6)$$

and compute PJR as the ratio of the announcement window CAR and the event window CAR,

$$PJR_i = \frac{CAR_i^{T-1, T+2}}{CAR_i^{T-21, T+2}}. \quad (7)$$

One notable implementation issue when computing PJR is that the denominator of the metric may occasionally be close to zero. To account for this issue, [Weller \(2018\)](#) drops the announcements for which the absolute event-window CAR is smaller than  $\sqrt{24}\sigma_i$ , where  $\sigma_i$  is the standard deviation of  $r_i$  over the preceding month. We do the same.

Table 10 examines PJRs around the switch to continuous trading and reports no evidence of price discovery improvements. The estimates show that PJRs in Taiwan are relatively close to one. As such, price discovery mainly occurs close to the announcements. More importantly, there appears to be no evidence of changes in PJR after the switch to continuous trading. The univariate results in Panel A and the regression specification 1 in Panel B focus on Taiwan only and report insignificant changes in PJR. Specification 2 in Panel B adds Korean earnings announcements as controls and also reports no changes in PJRs, neither in Taiwan (the  $Post \times TWSE$  coefficient) nor Korea (the  $Post$  coefficient). As such, investors appear to impound the same amount of information into prices during the auction and continuous regimes. Taken together, these results are consistent with the notion that the switch to continuous trading affects adverse selection and liquidity mainly through the latency arbitrage channel rather than the price discovery channel.

All things considered, the data paint a rich picture of changes in market participation. On the one hand, total volume increases (unconditionally) upon the switch to continuous trading. On the other hand, the additional volume appears to be generated in large part by latency arbitrageurs, while volume generated by (some of) the traditional market participants declines. These results point to a notable dichotomy in the interests of market operators and various investor groups. Total volume positively affects exchange revenues, which are derived to a great degree from market access fees paid by traders on a per-share basis. If the TWSE experience is indicative of

the tradeoffs between discrete and continuous trading in a typical market, it may not be in the interest of modern market operators to move to the batch auction structure; trading volume may decline as a result. This logic echoes [Budish, Lee, and Shim \(2020\)](#), who suggest that modern stock exchanges are unlikely to move away from the status quo of continuous trading without regulatory encouragement.

## 4.5 Robustness

For several key variables used in this study, we chose estimation horizons that are commonly used in the literature. Specifically, we rely on 30-second horizons when we estimate price impacts and use 60-second horizons for return autocorrelation and price delay metrics. In Panels A and B of [Table 11](#), we ask if our results are robust to alternative horizons. The data indicate that they are. In the DID regression specification that uses volume and volatility controls, all above-mentioned variables remain statistically significant and have similar economic magnitudes to those reported in the main tables.

[[Table 11](#)]

An issue that deserves additional examination is that of trade signing during the auction regime. In an earlier section, we suggest that the logic of the Lee-Ready algorithm is applicable in the auction environment. To examine the sensitivity of our main results to this algorithm choice, we examine two additional approaches to trade signing. The first approach assumes that price discovery is mainly driven by trades, and the second approach assumes that it is mostly driven by quotes. As such, the two approaches largely straddle the spectrum of trade signing possibilities.

When working with the first approach, we sign trades based on the price change that follows them, effectively assuming that every such trade is informed. That is, if a trade is followed by a price increase (decline), we consider it buyer-initiated (seller-initiated). In the Lee-Ready tradition, trades not followed by a price change are signed as the trade preceding them. In the second

approach, we sign trades against the quotes that follow them rather than against the quotes that precede them. For example, consider a trade at NTD 10.00 that follows the quotes with NTD 9.99 on the bid and 10.00 on the offer and results in new quotes with NTD 10.00 on the bid and 10.01 on the offer. The Lee-Ready algorithm would sign this trade as buyer-initiated, whereas the second approach discussed here signs it as a seller-initiated, effectively assuming that price discovery had occurred before the trade was executed.

Panel C of Table 11 shows that no matter how trades are signed, price impacts increase after the switch to continuous trading. As such, trade signing choices do not affect the qualitative nature of our main result. Quantitatively, the results are strongest when we assume that auction price discovery is driven mainly by quotes. This assumption is not without merit given that the auction environment reduces profitability of latency arbitrage strategies. Notably, our base case Lee-Ready results are much closer in magnitude to the alternative that assumes trade-driven price discovery. As such, it may be appropriate to think of the economic magnitudes of price impact changes reported in the main analysis as relatively conservative.

## **5. Conclusion**

Market structure theory suggests that the continuous limit order book – market design that dominates modern equity trading – is prone to generating adverse selection. For every market maker order that may be attempting to change a stale quote, 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, and even assuming equal speeds by all market participants, the odds of replacing a stale quote before it is picked off are relatively low. As such, the adverse selection cost born by market makers is high. To compensate for this cost, spreads are kept wider than they would be 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. The data support the above-mentioned theory predictions, in that adverse selection increases significantly. In addition, market maker inventory costs increase, consistent with the notion that latency arbitrage complicates inventory management. The total liquidity effect of the TWSE move is therefore negative; trading costs increase, and displayed liquidity declines.

The increase in trading costs is sufficiently large to affect market participation by some end-users of liquidity, that is, uninformed investors who come to the market for the purpose of traditional asset exchange rather than to engage in latency arbitrage. We find that trading volume generated by such investors declines substantially, potentially reducing gains from trade. Notably, this decline is more than compensated for by an increase in trading volume generated in the process of latency arbitrage. In turn, arbitrage activity comes with moderate improvements in price efficiency. Finally, the greater trading costs that follow the switch to continuous trading do not appear to materially affect investors who engage in fundamental price discovery. Proxied for by information incorporation into prices prior to earnings announcements, market participation by these investors does not change.

Our results offer new empirical evidence to 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 benefiting the end-users of liquidity. On the other hand, the continuous design comes with greater trading volumes boosted by arbitrage activity benefiting the exchanges. Given the high fixed costs of running an exchange, it is unlikely that market operators will willingly change the status quo, especially if the change may reduce trading volumes. In the meantime, it appears that for the continuous order book design to be welfare-improving, the end consumers of liquidity must heavily discount increases in

trading costs and put a substantial premium on moderate improvements in price efficiency.

## References

- Aït-Sahalia, Y., and M. Sağlam, 2017, “High frequency market making: Optimal quoting,” Working paper, Princeton University.
- Amihud, Y., H. Mendelson, and B. Lauterbach, 1997, “Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange,” *Journal of Financial Economics*, 45(3), 365–390.
- Aquilina, M., E. B. Budish, and P. O’Neill, 2020, “Quantifying the high-frequency trading “arms race”: A simple new methodology and estimates,” *Chicago Booth Research Paper*, (20-16).
- Barclay, M. J., T. Hendershott, and C. M. Jones, 2008, “Order consolidation, price efficiency, and extreme liquidity shocks,” *Journal of Financial and Quantitative Analysis*, pp. 93–121.
- Baron, M., J. Brogaard, B. Hagströmer, and A. Kirilenko, 2019, “Risk and return in high-frequency trading,” *Journal of Financial and Quantitative Analysis*, 54(3), 993–1024.
- Biais, B., T. Foucault, and S. Moinas, 2015, “Equilibrium fast trading,” *Journal of Financial economics*, 116(2), 292–313.
- 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.
- 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, 2020, “A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?,” Working paper, National Bureau of Economic Research.
- Campbell, J. Y., T. Ramadorai, and A. Schwartz, 2009, “Caught on tape: Institutional trading, stock returns, and earnings announcements,” *Journal of Financial Economics*, 92(1), 66–91.
- 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.
- Comerton-Forde, C., and J. Rydge, 2006, “The influence of call auction algorithm rules on market efficiency,” *Journal of Financial Markets*, 9(2), 199–222.
- 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.
- Ellul, A., H. S. Shin, and I. Tonks, 2005, “Opening and closing the market: Evidence from the London Stock Exchange,” *Journal of Financial and Quantitative Analysis*, pp. 779–801.
- Foucault, T., R. Kozhan, and W. W. Tham, 2017, “Toxic arbitrage,” *Review of Financial Studies*, 30(4), 1053–1094.
- Foucault, T., and S. Moinas, 2019, “Is Trading Fast Dangerous?,” *Global Algorithmic Capital Markets: High Frequency Trading, Dark Pools, and Regulatory Challenges*, pp. 9–28.

- Haas, M., M. Khapko, and M. Zoican, 2020, “Speed and learning in high-frequency auctions,” *Journal of Financial Markets*, 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 P. C. Moulton, 2011, “Automation, speed, and stock market quality: The NYSE’s hybrid,” *Journal of Financial Markets*, 14(4), 568–604.
- Henke, H., and B. Lauterbach, 2005, “Firm-initiated and exchange-initiated transfers to continuous trading: Evidence from the Warsaw Stock Exchange,” *Journal of Financial Markets*, 8(3), 309–323.
- 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.
- Kyle, A. S., 1985, “Continuous Auctions and Insider Trading,” *Econometrica*, 53(6), 1315–1335.
- Lee, C., and M. Ready, 1991, “Inferring trade direction from intraday data,” *Journal of Finance*, 46(2), 733–746.
- Madhavan, A., 1992, “Trading mechanisms in securities markets,” *Journal of Finance*, 47(2), 607–641.
- Muscarella, C. J., and M. S. Piwowar, 2001, “Market microstructure and securities values:: Evidence from the Paris Bourse,” *Journal of Financial Markets*, 4(3), 209–229.

- Pagano, M. S., L. Peng, and R. A. Schwartz, 2013, “A call auction’s impact on price formation and order routing: Evidence from the NASDAQ stock market,” *Journal of Financial Markets*, 16(2), 331–361.
- Pagano, M. S., and R. A. Schwartz, 2003, “A closing call’s impact on market quality at Euronext Paris,” *Journal of Financial Economics*, 68(3), 439–484.
- Riccò, R., and K. Wang, 2020, “Frequent Batch Auctions vs. Continuous Trading: Evidence from Taiwan,” Working paper, NHH Norwegian School of Economics.
- Roll, R., 1984, “A simple implicit measure of the effective bid-ask spread in an efficient market,” *Journal of Finance*, 39(4), 1127–1139.
- Roşu, I., 2019, “Fast and slow informed trading,” *Journal of Financial Markets*, 43, 1–30.
- Schnitzlein, C. R., 1996, “Call and continuous trading mechanisms under asymmetric information: An experimental investigation,” *Journal of Finance*, 51(2), 613–636.
- 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.
- Theissen, E., 2000, “Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets,” *Journal of Financial Markets*, 3(4), 333–363.
- Weller, B. M., 2018, “Does algorithmic trading reduce information acquisition?,” *Review of Financial Studies*, 31(6), 2184–2226.

**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 midpoints scaled by the average midpoint. Quote midpoint is the average between the TWSE best bid and best offer prices.

	Mean	Median	Std. Dev.	10th	90th
Market cap., NTD million	282,447	118,614	840,716	54,028	474,412
Price, NTD	181.92	67.65	491.58	14.73	372.05
Number of trades	1,076	972	622	305	1,920
Volume, share thousand	9,550	4,840	16,294	553	23,116
Volatility, bps.	1.43	1.23	0.83	0.52	2.74

**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 quote midpoint 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 quote midpoint 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 quote midpoint immediately preceding the trade. *Price impact* is twice the signed difference between the quote midpoint immediately preceding the trade and the midpoint 30 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 quote midpoints. 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 midpoint 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  $R^2$ s 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 3 contains estimation details). The two  $R^2$ s 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.	23.41	20.44	10.94	12.09	39.93
Quoted depth, share thousand	447.5	92.7	928.2	8.1	943.9
Effective spread, bps.	19.12	15.63	9.33	10.16	33.89
Price impact, bps.	10.84	9.62	5.46	5.09	19.00
Realized spread, bps.	8.27	6.10	8.15	0.04	18.74
Panel B: Price efficiency metrics					
Return autocorrelation (60s)	0.11	0.11	0.02	0.08	0.14
Return autocorrelation (PC1)	0.33	0.33	0.09	0.23	0.42
Price delay (60s)	0.08	0.06	0.08	0.00	0.19
Price delay (PC1)	0.77	0.86	0.04	0.81	0.89

**Table 4**  
**Adverse Selection**

The table examines changes in adverse selection of liquidity providers (proxied by the price impacts) 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:

$$price\ impact_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_i + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \epsilon_{it},$$

where *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*; and *Volatility* is the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard errors are in parentheses. \*\*\* indicates statistical significance at the 1% level.

	[1]		[2]	
<b>Panel A: Univariate results</b>				
Pre	10.84			
Post	13.78	***		
<b>Panel B: Regression results</b>				
<i>Post</i>	0.010 (0.04)		-0.074 (0.02)	***
<i>TWSE</i>	-0.240 (0.04)	***	-0.121 (0.02)	***
<i>Post</i> × <i>TWSE</i>	0.460 (0.05)	***	0.235 (0.03)	***
<i>Volume</i>			-0.038 (0.02)	***
<i>Volatility</i>			0.570 (0.02)	***
<i>Intercept</i>	-0.002 (0.04)		0.033 (0.03)	
Adj. R <sup>2</sup>	0.028		0.310	
Obs.	24,144		24,144	
<b>Panel C: Regression: alternative sample periods</b>				
<i>Post</i> × <i>TWSE: excluding March</i>	0.320 (0.06)	***	0.143 (0.04)	***
<i>Post</i> × <i>TWSE: full sample</i>	0.280 (0.05)	***	0.177 (0.04)	***

**Table 5**  
**Displayed Liquidity**

The table examines changes in quoted spread and depth 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:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \varepsilon_{it},$$

where  $DepVar$  is the quoted spread or quoted depth,  $Post$  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  $t$ ; and  $Volatility$  is the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard deviations are in parentheses. \*\*\* and \*\* indicate statistical significance at the 1% and 5% levels.

	Quoted spread				Quoted depth			
	[1]		[2]		[3]		[4]	
Panel A: Univariate results								
Pre	23.41				447.5			
Post	25.74	***			322.4	***		
Panel B: Regression results								
<i>Post</i>	-0.076	**	-0.045		0.224	***	0.153	***
	(0.04)		(0.04)		(0.05)		(0.04)	
<i>TWSE</i>	-0.468	***	-0.471	***	0.190	***	0.194	***
	(0.03)		(0.03)		(0.04)		(0.03)	
<i>Post</i> × <i>TWSE</i>	0.901	***	0.907	***	-0.375	***	-0.380	***
	(0.06)		(0.05)		(0.06)		(0.05)	
<i>Volume</i>			-0.267	***			0.636	***
			(0.02)				(0.02)	
<i>Volatility</i>			0.178	***			-0.444	***
			(0.02)				(0.01)	
<i>Intercept</i>	0.070		0.042		-0.016	***	-0.096	***
	(0.05)		(0.04)		(0.04)		(0.03)	
Adj. R <sup>2</sup>	0.086		0.119		0.01		0.201	
Obs.	24,144		24,144		24,144		24,144	
Panel C: Regression: alternative sample periods								
<i>Post</i> × <i>TWSE: excluding March</i>	0.799	***	0.820	***	-0.358	***	-0.419	***
	(0.05)		(0.04)		(0.06)		(0.04)	
<i>Post</i> × <i>TWSE: full sample</i>	0.751	***	0.703	***	-0.363	***	-0.323	***
	(0.04)		(0.03)		(0.05)		(0.04)	

**Table 6**  
**Trading Costs**

The table examines changes in effective 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 DID regression of the following form:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \varepsilon_{it},$$

where  $DepVar$  is the effective 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  $t$ ; and  $Volatility$  is the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard deviations are in parentheses. \*\*\* and \*\* indicate statistical significance at the 1% and 5% levels.

	Effective spread				Realized spread			
	[1]		[2]		[3]		[4]	
Panel A: Univariate results								
Pre	19.12				8.27			
Post	23.01		***		9.23		***	
Panel B: Regression results								
<i>Post</i>	-0.095	***	-0.087	**	-0.187	***	-0.100	***
	(0.04)		(0.04)		(0.03)		(0.02)	
<i>TWSE</i>	-0.609	***	-0.596	***	-0.172	***	-0.281	***
	(0.03)		(0.03)		(0.03)		(0.02)	
<i>Post</i> × <i>TWSE</i>	1.175	***	1.149	***	0.341	***	0.545	***
	(0.05)		(0.04)		(0.04)		(0.04)	
<i>Volume</i>			-0.160	***			-0.066	***
			(0.01)				(0.01)	
<i>Volatility</i>			0.176	***			-0.442	***
			(0.02)				(0.02)	
<i>Intercept</i>	0.045		0.033		0.089	***	0.047	
	(0.04)		(0.04)		(0.03)		(0.03)	
Adj. R <sup>2</sup>	0.147		0.162		0.007		0.238	
Obs.	24,144		24,144		24,144		24,144	
Panel C: Regression: alternative sample periods								
<i>Post</i> × <i>TWSE: excluding March</i>	1.083	***	1.072	***	0.422	***	0.595	***
	(0.04)		(0.04)		(0.04)		(0.03)	
<i>Post</i> × <i>TWSE: full sample</i>	0.933	***	0.885	***	0.491	***	0.562	***
	(0.04)		(0.03)		(0.04)		(0.03)	

**Table 7**  
**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 univariate results. Panels B and C report results from a DID regression of the following form:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \varepsilon_{it},$$

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*; and *Volatility* is the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard deviations are in parentheses. \*\*\* indicate statistical significance at the 1% level.

	Return autocorrelation				Price delay			
	60s		PC1		60s		PC1	
	[1]		[2]		[3]		[4]	
Panel A: Univariate results								
Pre	0.112		0.329		0.850		0.767	
Post	0.095	***	0.287	***	0.704	***	0.684	***
Panel B: Regression results								
<i>Post</i>	0.030		0.023	***	-0.329	***	-0.076	***
	(0.02)		(0.01)		(0.04)		(0.01)	
<i>TWSE</i>	0.044		-0.032	***	0.158	***	0.029	***
	(0.02)		(0.01)		(0.03)		(0.01)	
<i>Post</i> × <i>TWSE</i>	-0.181	***	-0.056	***	-0.215	***	-0.058	***
	(0.03)		(0.01)		(0.06)		(0.02)	
<i>Volume</i>	-0.036	***	0.002		0.074	***	0.020	***
	(0.01)		(0.00)		(0.01)		(0.00)	
<i>Volatility</i>	-0.083	***	-0.027	***	-0.112	***	-0.032	***
	(0.01)		(0.00)		(0.02)		(0.01)	
<i>Intercept</i>	0.011		0.351	***	0.287	***	0.806	***
	(0.02)		(0.00)		(0.04)		(0.01)	
Adj. R <sup>2</sup>	0.014		0.038		0.069		0.086	
Obs.	24,144		23,353		23,149		23,950	
Panel C: Regression: alternative sample periods								
<i>Post</i> × <i>TWSE: excluding March</i>	-0.179	***	-0.050	***	-0.154	***	-0.031	
	(0.03)		(0.01)		(0.06)		(0.02)	
<i>Post</i> × <i>TWSE: full sample</i>	-0.184	***	-0.051	***	-0.051		-0.010	
	(0.03)		(0.01)		(0.06)		(0.02)	

**Table 8**  
**Volatility and Volume**

The table examines changes in volume and volatility 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. Panels B and C report the results of a pooled DID regression of the following form:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \varepsilon_{it},$$

where *DepVar* is trading volume or volatility, *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*; and *Volatility* the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard errors are in parentheses. \*\*\* indicates statistical significance at the 1% level.

	Volatility		Volume	
	[1]	[2]	[3]	[4]
Panel A: Univariate results				
Pre	1.43		9,551	
Post	1.96	***	11,795	***
Panel B: Regression results				
<i>Post</i>	0.116	***	-0.125	***
	(0.04)		(0.03)	
<i>TWSE</i>	-0.160	***	-0.113	***
	(0.03)		(0.02)	
<i>Post</i> × <i>TWSE</i>	0.290	***	0.175	***
	(0.05)		(0.04)	
<i>Volume</i>			0.557	***
			(0.01)	
<i>Volatility</i>				0.879
				(0.02)
<i>Intercept</i>	-0.285	***	-0.052	***
	(0.04)		(0.04)	
Adj. R <sup>2</sup>	0.036		0.508	
Obs.	24,144		24,144	
Panel C: Regression: alternative sample periods				
<i>Post</i> × <i>TWSE</i> : <i>excluding March</i>	0.240	***	0.123	***
	(0.05)		(0.04)	
<i>Post</i> × <i>TWSE</i> : <i>full sample</i>	0.170	***	0.118	***
	(0.06)		(0.04)	
			0.200	***
			(0.07)	
			0.080	***
			(0.07)	
			0.009	***
			(0.05)	
			-0.036	***
			(0.05)	

**Table 9**  
**Low-Toxicity Volume**

The table examines changes in low-toxicity volume around the move to continuous trading. We define such volume as originating from trades, whose price impacts are negative and use low-toxicity volume to proxy for market participation by the uninformed users of liquidity (the end-users). 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:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volatility_{it} + \varepsilon_{it},$$

where *DepVar* is low-toxicity 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, and *Volatility* the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard errors are in parentheses. \*\*\* and \*\* indicate statistical significance at the 1% and 5% levels.

	[1]		[2]	
<b>Panel A: Univariate results</b>				
Pre	1,179			
Post	392	***		
<b>Panel B: Regression results</b>				
<i>Post</i>	0.239	***	0.014	***
	(-0.03)		(0.03)	
<i>TWSE</i>	0.430	***	0.565	**
	(0.04)		(0.03)	
<i>Post</i> × <i>TWSE</i>	-0.840	***	-1.092	***
	(0.05)		(0.04)	
<i>Volatility</i>			0.603	***
			(0.02)	
<i>Intercept</i>	-0.147	***	-0.104	***
	(0.03)		(0.03)	
Adj. R <sup>2</sup>	0.053		0.400	
Obs.	24144		24,144	
<b>Panel C: Regression: alternative sample periods</b>				
<i>Post</i> × <i>TWSE</i> : <i>excluding March</i>	-0.079	***	-0.988	***
	(0.04)		(0.04)	
<i>Post</i> × <i>TWSE</i> : <i>full sample</i>	-0.840	***	-0.941	***
	(0.07)		(0.06)	

**Table 10**  
**Informed Trading Around Earnings Announcements**

The table examines changes in the way earnings news are incorporated into prices around the move to continuous trading. The main metric is the price jump ratio,  $PJR$ , computed as return immediately surrounding an earnings announcement divided by return that includes three weeks preceding the announcement,

$$PJR_i = \frac{CAR_i^{T-1, T+2}}{CAR_i^{T-21, T+2}},$$

where  $CAR_i^{T-1, T+2}$  is the cumulative market-adjusted return for the announcement  $i$  from day  $T - 1$  to day  $T + 2$ , with  $T$  being the announcement date, and  $CAR_i^{T-21, T+2}$  is the same metric computed from day  $T - 21$  to day  $T + 2$ . The intuition is based on [Weller \(2018\)](#), who posits that informed investors will impound earnings information into prices prior to the announcement if the trading environment is conducive to such activity. The treatment sample consists of 100 largest TWSE stocks, and the control group is 100 matched KRX stocks. We caution that not all sample stocks have earnings announcements during the sample period, and as such this test is not as balanced as those reported in the previous tables. The sample period spans November 1, 2019 to July 30, 2020. Panel A contains univariate results. Panel B reports the results of a pooled DID regression of the following form:

$$PJR_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_i + \beta_3 Post_t \times TWSE_{it} + \varepsilon_{it},$$

where  $PJR$  is the jump ratio,  $Post$  is an indicator variable that equals to 1 in the post-switch period and zero otherwise, and  $TWSE$  is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks. White-robust standard errors are in parentheses. \*\*\* indicates statistical significance at the 1% level.

	TWSE only	TWSE & KRX		
	[1]	[2]		
<b>Panel A: Univariate results</b>				
Pre	0.91			
Post	0.93			
<b>Panel B: Regression results</b>				
<i>Post</i>	0.028 (0.23)	-0.374 (0.25)		
<i>TWSE</i>		-0.011 (0.26)		
<i>Post</i> × <i>TWSE</i>		0.402 (0.34)		
<i>Intercept</i>	0.907 (0.12)	*** 0.918 (0.23)	***	***
Adj. R <sup>2</sup>	0.010	0.060		
Obs.	181	239		

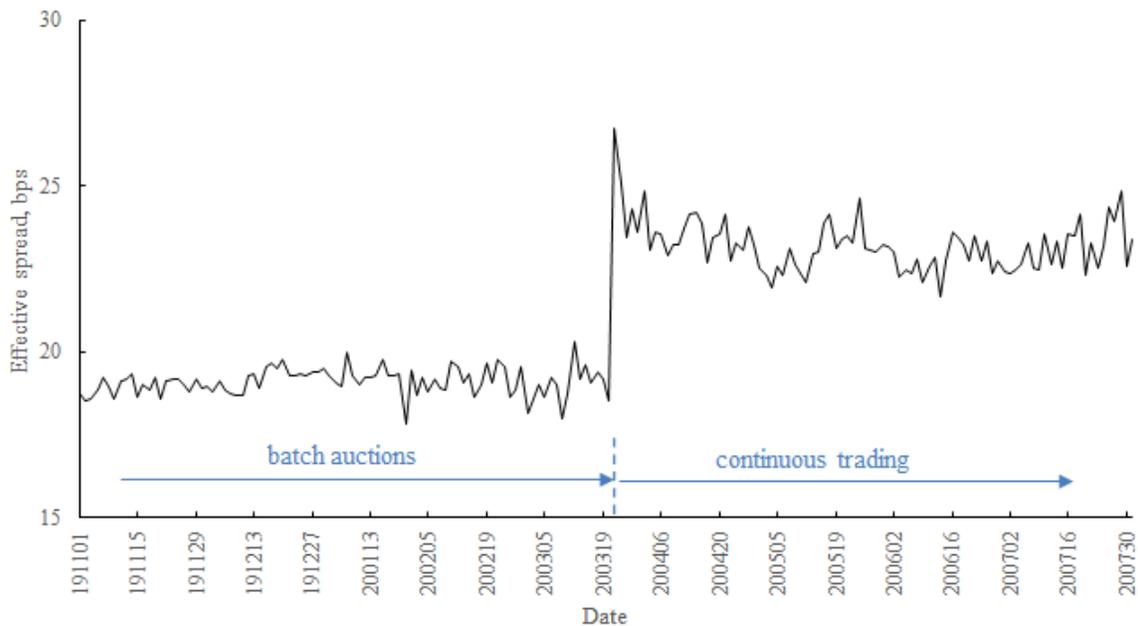
**Table 11**  
**Robustness**

The table contains regression results for price impacts and price efficiency metrics estimated at various horizons and across the spectrum of trade classification. For price impacts, we use 10, 15, 60, and 300-second horizons (Panel A) and for the price efficiency metrics – 10, 30, and 300-second horizons (Panel B). To examine the trade classification spectrum, in Panel C we report two opposite approaches for trade signing that assume that price discovery is driven either mainly by trades or mainly by quotes. The former approach signs trades based on the midquote change that follows them. That is, trades that are followed by price increases (declines) are considered buyer-initiated (seller-initiated). In turn, the latter approach uses post-trade quotes instead of the pre-trade quotes to sign trades, effectively assuming that quotes lead price discovery. For comparison to the findings reported throughout the paper, we also report the result from the Lee-Ready classification (the base case). 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, the pre-event window includes November 2019 through January 2020, and the post-event window includes May through July 2020. The table reports the coefficient estimates on the  $Post_t \times TWSE_{it}$  variable from a DID regression of the following form:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 TWSE_{it} + \beta_3 Post_t \times TWSE_{it} + \delta_1 Volume_{it} + \delta_2 Volatility_{it} + \epsilon_{it},$$

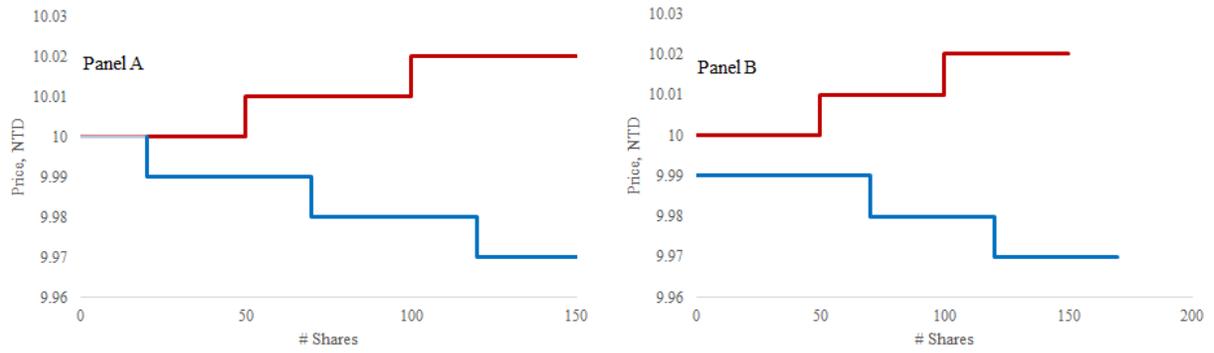
where  $DepVar$  are the price impact, autocorrelation, and price delay metrics,  $Post$  is an indicator variable that equals to 1 for the post-event period and zero otherwise;  $TWSE$  is an indicator variable that equals to 1 for the TWSE stocks and 0 for the KRX stocks;  $Volume$  is daily trading volume in stock  $i$  on day  $t$ ; and  $Volatility$  is the difference between the highest and lowest midpoints scaled by the average midpoint. 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. White-robust standard deviations are in parentheses. \*\*\* indicate statistical significance at the 1% level.

Panel A: Price impact horizons								
	10 seconds		15 seconds		60 seconds		300 seconds	
Price impact	0.420	***	0.304	***	0.269	***	0.313	***
	(0.04)		(0.03)		(0.03)		(0.04)	
Panel B: Price efficiency								
	10 seconds		30 seconds		300 seconds			
Autocorrelation	-0.093	***	-0.222	***	-0.095	***		
	(0.04)		(0.03)		(0.04)			
Price delay	-0.161	***	-0.217	***	-0.242	***		
	(0.05)		(0.08)		(0.08)			
Panel C: Price impacts across the trade classification spectrum								
	Price discovery occurs mainly				Base case			
	via trades		via quotes		(Lee-Ready)			
Price impact	0.159	***	1.386	***	0.235	***		
	(0.04)		(0.03)		(0.03)			



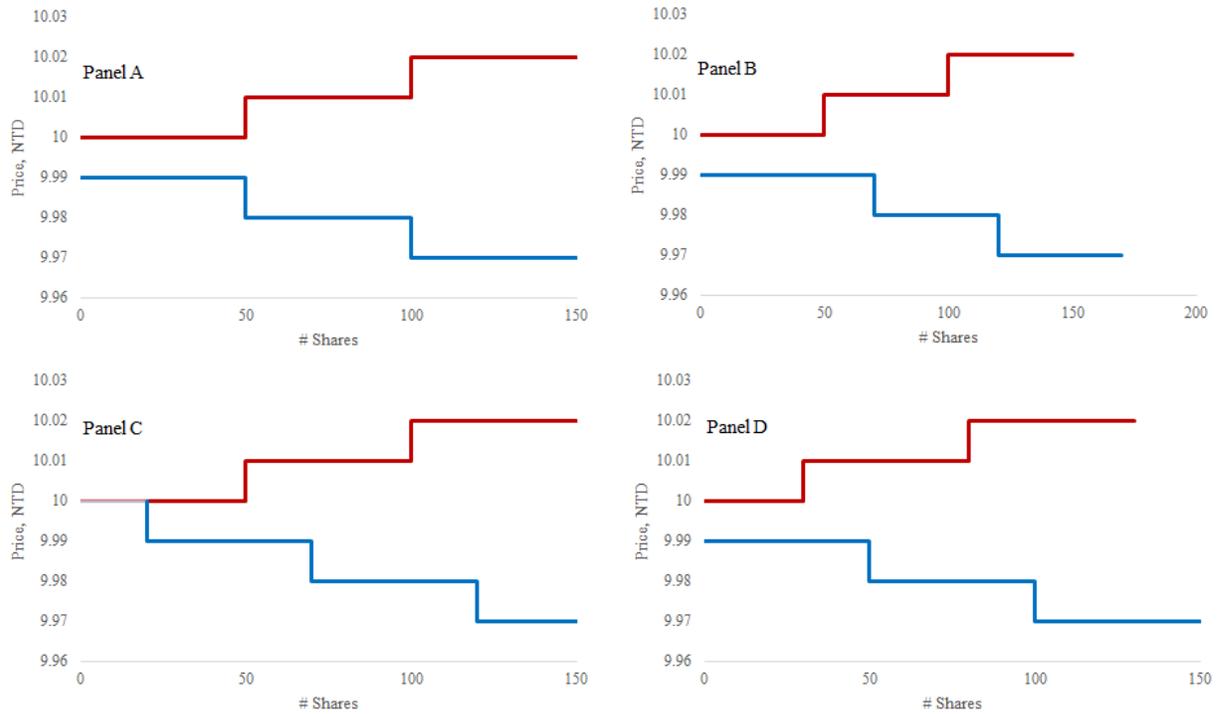
**Figure 1**  
**Trading costs around the switch to continuous trading**

The figure plots the effective spreads, our proxy for trading costs, from November 2019 through July 2020. The sample consists of 100 largest TWSE stocks. Effective spread is the signed difference between the trade price and the corresponding quote midpoint, scaled by the midpoint. We use the Lee and Ready (1991) algorithm to sign trades. In Section 3, we discuss assumptions required to compute effective spreads in the auction environment. For aggregation, effective spreads are first volume-weighted within each stock-day and then averaged across stocks for each day.



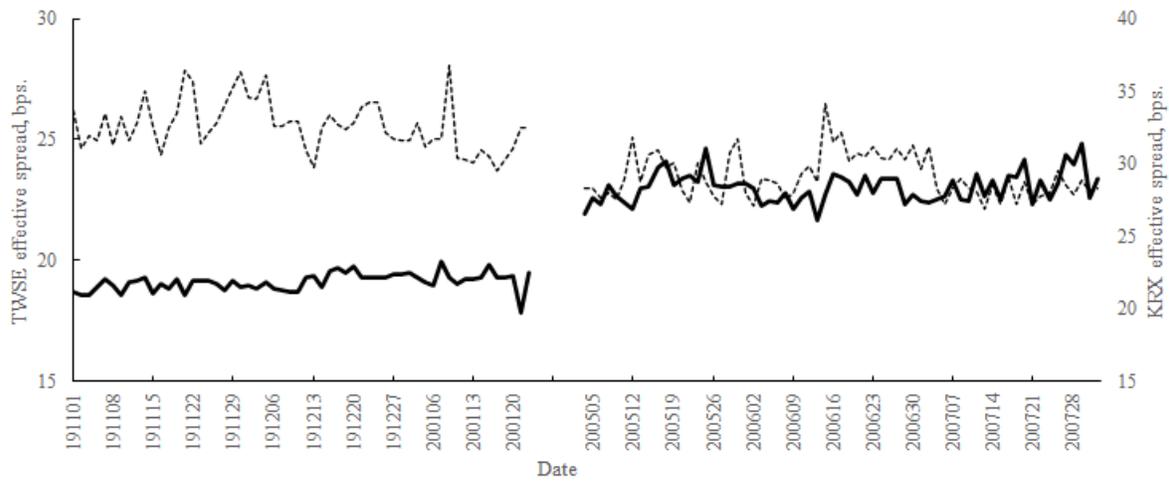
**Figure 2**  
**An auction market example**

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 3**  
**A continuous market example**

The figure plots an example of trading in a continuous market for comparison with the auction market example in Figure 2. 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.



**Figure 4**

**TWSE and KRX trading costs during the main sample period**

The figure plots the TWSE effective spreads (solid line) and the KRX effective spreads (dashed line) during the main event window spanning November and December 2019 and January 2020 in the pre-event period and May-July 2020 in the post-event period. The sample consists of 100 largest TWSE stocks and their KRX matches. Effective spread is the signed difference between the trade price and the corresponding quote midpoint, scaled by the midpoint. As is conventional, we drop trading days when either the TWSE or KRX is closed for a holiday. Given the Lunar New Year celebrations in late January 2020 and several other holidays, the pre-event window is four days shorter than the post-event window.