

Discussion Papers In Economics And Business

Effectiveness of Trading Pauses: Evidence from the Tokyo Stock Exchange

Akitada Kasahara Masahiro Yamada

Discussion Paper 26-03-Rev.

October 2026

Graduate School of Economics
The University of Osaka, Toyonaka, Osaka 560-0043, JAPAN

Effectiveness of Trading Pauses: Evidence from the Tokyo Stock Exchange*

Akitada Kasahara† Masahiro Yamada‡

Abstract

We estimate the causal effect of single-stock trading pauses on market quality using tick-by-tick order book data from the Tokyo Stock Exchange (TSE). Using percentage distances to yen-denominated triggering thresholds as an instrument, we find that pauses extending the opening auction reduce subsequent volatility by roughly a quarter, consistent with reduced transitory price variation. Order flows suggest that extra auction time lets aggressively priced limit orders offset the initial imbalance. Reductions are larger for stocks with lower turnover and smaller after overnight price declines. We detect no robust volatility reduction for pauses interrupting continuous trading.

Keywords: Trading pauses, Circuit breakers, Volatility, Liquidity, Price discovery

JEL Classification: G12, G14, G18

*We thank Dimitris Papanikolaou (the editor) and an anonymous referee for helpful comments. We are also grateful to Wataru Ohta and seminar participants at SWET 2024, the 6th Annual Autumn Meeting of the Nippon Finance Association, and the Three Asian Countries Finance Conference 2025 for their valuable feedback, and to the Japan Exchange Group for generously providing the data. All remaining errors are our own. This research was supported by JSPS KAKENHI Grant Number JP22K13428 and JP25K00650.

†Corresponding author. Graduate School of Economics, the University of Osaka. Email: akitadak@econ.osaka-u.ac.jp.

‡School of Management, Tokyo University of Science. Email: m.yamada@rs.tus.ac.jp

1 Introduction

Trading pauses—temporary suspensions of execution triggered by sharp price movements—are among the most widely adopted regulatory tools in modern financial markets. Despite their prevalence at both the market-wide and individual-stock levels, their effectiveness remains actively debated, as the theoretical literature yields ambiguous predictions and credible causal evidence is scarce (Harris, 1997). Because trading pauses are triggered precisely when volatility is elevated, naive comparisons between paused and non-paused episodes are confounded by selection.

We estimate the causal effect of single-stock trading pauses on market quality using tick-by-tick order book data from the Tokyo Stock Exchange (TSE). Most pauses on the TSE do not interrupt continuous trading. In our main sample of the 500 largest securities among First Section stocks and J-REITs from 2014 to 2021, 86% of pauses occur at the 9:00 AM open, before continuous trading in the stock begins. Such a pause delays the opening call auction while orders can still be submitted; the range of prices at which the auction may execute widens every three minutes until the order imbalance can be cleared, and the stock then opens through the same batch auction that would have opened it without a pause. These pauses therefore raise a different question: does giving the opening auction more time improve its outcome? This question is closer to the analysis of batch auctions by Madhavan (1992) than to the circuit-breaker debate as usually framed. We provide causal evidence on these opening-auction extensions, together with order-level evidence on how liquidity is supplied during the added auction time, and we study the rarer pauses that interrupt continuous trading separately.

For the opening analysis, our instrumental variable (IV) is the percentage distance from the opening reference price, essentially the previous day's close, to the pause-triggering threshold. Because the thresholds follow a fixed yen-denominated schedule, this distance varies across stocks with different price levels and is determined before the day's orders arrive. This variation identifies the effect of a pause within a single regulatory environment, without drawing the counterfactual from a period or market without pauses.

Our empirical analysis yields four sets of findings.

1. **Opening-auction extensions reduce volatility.** An extension reduces realized volatility over the first 30 minutes of trading by roughly a quarter, consistent with reduced transitory price variation. Quoted spreads narrow, and trading activity does not decline significantly. Among paused stocks, the opening price that emerges from the extended auction also tracks the day's closing price more closely than the price at which the auction would have cleared at 9:00 AM.
2. **Counter-orders are placed ahead of the widening range.** During opening-auction extensions, new orders against the initial imbalance account for more than half of classifiable new-order volume. Against-direction limit orders arrive throughout each complete three-minute cycle rather than only just before the next widening, and most are priced beyond the current edge of the range, so the widening range moves toward liquidity that is already waiting. The widening times are known in advance, which may help liquidity providers plan their submissions, although we do not separately identify the role of this predictability.
3. **Volatility reductions vary with turnover and market conditions.** The reduction is present among large and highly liquid stocks and is larger for stocks with lower turnover. It is smaller after the stock's overnight price declines, after large market-wide overnight moves, and for stocks that have recently been volatile or falling; after earnings announcements, its point estimate is close to zero. Unlike the volatility reduction, the narrowing of quoted spreads is not observed among the largest stocks, those in the TOPIX 100.
4. **Pauses that interrupt continuous trading show no robust volatility reduction.** For pauses triggered during continuous trading, which mostly occur in smaller stocks, we detect no robust change in volatility after the pause. The estimate for price-driven pauses, which interrupt an active market, is close to zero and imprecise, while the evidence of a reduction for order-book-evaporation pauses, which begin when no price can match buy and sell orders, is marginal; neither estimate is precise enough to establish a difference from the effect at the open. OLS estimates on the same

samples are positive, as in prior work, whereas the IV estimates are close to zero or negative, consistent with selection contributing to the positive association.

Taken together, the evidence of a volatility reduction comes from pauses that extend a call auction, at the morning open and, less precisely, at the 12:30 PM open, consistent with the view that batch auctions aggregate information more reliably when given time (Madhavan, 1992). For pauses that interrupt continuous trading, our estimates do not robustly establish either a benefit or a harm.

We also examine two concerns that the circuit-breaker literature associates with pauses. First, pausing some stocks at the open could impair price discovery on common factors and thereby harm other stocks in the same industry. Our IV estimates detect no statistically significant adverse spillovers, but their confidence sets allow economically meaningful increases in volatility. Second, the anticipation of a pause could draw prices toward the trigger, the magnet effect of Subrahmanyam (1994). The opening auction provides no preceding period of continuous trading in which to test this channel, so we test it for pauses during continuous trading, using variation generated by the same yen-denominated range-step schedule, and find no evidence consistent with a magnet effect.

The analysis of continuous-session pauses also brings our setting closer to the U.S. trading pauses studied by Hautsch and Horvath (2019) (described in the literature review below). The U.S. pauses do not apply to the opening auction, and a large number of them occur in securities outside the S&P 500 and Russell 1000 (Moise and Flaherty, 2017). Pauses at the TSE open, by contrast, are common among large stocks: about 9% of stock–dates in our top-500 sample involve one, which gives our design a sizable sample of pauses in large and liquid stocks. Differences in event populations, institutional details, and empirical methods nonetheless limit direct comparison with the U.S. evidence.

Related Literature. Our study connects the literature on trading pauses to work on opening auctions and pre-opening price discovery. Madhavan (1992) shows that periodic batch auctions can aggregate information more efficiently and are more robust to information asymmetry, enabling markets to operate when continuous trading mechanisms break down. Biais et al. (1999) show that pre-opening indicative prices on the Paris Bourse become increasingly informative as the open approaches, and the design of opening and closing auctions has measurable consequences for price formation (Madhavan and Panchapagesan, 2000; Pagano and Schwartz, 2003). We provide causal evidence on a rule that operates within the auction: extending its duration in response to an order imbalance.

The theoretical literature on trading pauses yields ambiguous predictions. Proponents argue that pauses allow information aggregation, reduce panic, and protect liquidity providers (Greenwald and Stein, 1988, 1991; Kodres and O'Brien, 1994). Critics counter that pauses delay price discovery and merely postpone volatility, since the underlying trading pressures remain unresolved during the interruption (Kyle, 1988; Fama, 1989), and Subrahmanyam (1994) shows that the anticipation of a halt can increase volatility beforehand; Cho et al. (2003) provide evidence of this magnet effect for daily price limits, and Chen et al. (2024) show that market-wide circuit breakers can generate it in a dynamic equilibrium model.

As Harris (1997) emphasizes, a central challenge in empirical research on circuit breakers is to identify appropriate counterfactuals. Lee et al. (1994) observe higher trading volume and volatility following individual-stock trading halts on the NYSE, comparing halts to pseudo-halts matched by time, duration, and net-of-market returns. Corwin and Lipson (2000) find that order submissions and cancellations increase significantly during NYSE halts, but acknowledge the absence of a suitable counterfactual.

More recent studies examine rule-based trading pauses. Hautsch and Horvath (2019) compare Nasdaq trading pauses from 2010 to 2014 with matched price-change episodes from before the pause rules, which were the single-stock circuit breaker and, from 2013, the Limit Up–Limit Down (LULD) Plan; they find higher volatility and wider spreads after pauses, and higher volatility and trading volume beforehand. Other studies similarly construct control samples from periods or markets without pauses; for example, Jain et al. (2020) compare TSE trading pauses (special quotes) and LULD events with simulated pseudo-halts and find that neither delays price discovery, with mixed effects on volatility and a magnet effect for TSE pauses only at the upper edge of the price-limit range. Such comparisons face

the concern that unobservable factors may change between periods and that participants may behave differently when a pause mechanism is anticipated (Subrahmanyam, 1994).

Major European venues also implement single-stock pauses, known as volatility interruptions (Zimmermann, 2014). Like TSE pauses, they switch trading in the stock to a temporary call auction rather than halting it. Using data from the Spanish exchange, Abad and Pascual (2010) show that market conditions remain unstable after such interruptions but normalize within about 90 minutes, and Abad and Pascual (2007) find no magnet effect around them.

Finally, the TSE's price-setting mechanisms have been studied extensively. Hamao and Hasbrouck (1995) show that the TSE provides high immediacy through its public limit order system and document the operation of its trading-pause mechanism, Kim and Rhee (1997) show that daily price limits cause volatility spillovers, and Deb et al. (2017) find that they curb transitory volatility. Ohta (2011) and Xiao and Yamamoto (2020) study price discovery and order submission in the TSE's pre-opening period, and Bellia et al. (2020) examine the role of low-latency traders in the pre-opening and opening periods; none of these studies focuses on trading pauses.

The remainder of the paper is organized as follows. Section 2 describes the institutional background, including the auction-extension character of TSE pauses and the three-minute widening schedule. Section 3 describes the data and presents preliminary analyses. Section 4 presents the IV estimates of the effects of opening-auction extensions on volatility, liquidity, and price efficiency. Section 5 examines which component of volatility falls, how liquidity is supplied during the extension, how the effects vary across stocks, and whether pauses spill over to other stocks. Section 6 studies pauses that interrupt continuous trading, including tests for the magnet effect. Section 7 concludes.

2 Institutional Background

Abad and Pascual (2013) categorize price control mechanisms according to various criteria, such as whether their activation is rule-based or discretionary, and whether they apply market-wide or to individual securities. At the individual-stock level, the TSE employs a combination of trading pauses, signaled by special quotes (*tokubetsu-kehai*), and daily price limits as its primary price control mechanisms. Both are rule-based: they specify objective conditions under which executions are temporarily suspended or price movements are constrained, with the aim of mitigating excessive volatility and maintaining orderly markets during periods of stress.

The TSE's trading day consists of two sessions, each of which opens with a call auction.[1] Trading pauses can be triggered at any time, including during these opening auctions. As Section 3 documents, 86% of the pauses in our main sample occur at the 9:00 AM open; the remainder occur at the 12:30 PM open or during continuous trading.

How a pause works. A pause is triggered when the next trade would execute outside the *price-limit range*, which extends one range step L above and below the reference price. During continuous trading, the reference price is generally the last execution price; at the morning open, it is the exchange's base price, normally the previous day's final price or final quote, adjusted for dividends and stock splits. The base price is fixed for the day, whereas the reference price during continuous trading moves with each execution. The range step L is a fixed yen amount that depends on the level of the reference price according to the exchange's range-step schedule (Table 1); for 90% of stock–dates in our sample, it amounts to 1.5% to 2.7% of the reference price (Table 4). When a pause is triggered, the exchange does not execute the trade and instead displays a *special quote*: a bid (for a buy imbalance) or an offer (for a sell imbalance) one range step away from the reference price. This quote signals the imbalance to market participants. Throughout the paper, we use "trading pause" for the mechanism and "special quote" for

[1] During our 2014–2021 sample period, the morning session opened with a call auction at 9:00 AM, in which orders accumulated before the open were matched to set the opening price, and ran until 11:30 AM. The afternoon session opened with a call auction at 12:30 PM and ran until 3:00 PM, ending with a closing auction.

this displayed quote.[2]

The special quote is not a limit on order prices: during the pause, the order book remains open, and limit orders can be submitted, amended, or cancelled, including at prices beyond the special quote, subject to the daily price limit and standard order rules. We refer to the price at which the auction would clear on the orders currently in the book as the *matching price*. If counter-orders arrive and the auction can clear within the price-limit range, the stock trades through a call auction and the pause ends. Otherwise, the exchange moves the special quote every three minutes by the range step that applies at the current quote, widening the price-limit range toward the matching price, until the auction can clear. We refer to these scheduled moves as *updates*. Updates stop at the daily price limit; if the auction still cannot clear, the stock may remain untraded until the close.[3]

Pauses at the open and during continuous trading. Because the order book remains open during a pause and executions take place through a call auction, the economic character of a pause depends on when it is triggered. A pause at the 9:00 AM open occurs before any trading in the stock that day: it delays the opening auction and gives the auction more time to attract orders, but the stock still opens through the same auction that would have opened it without a pause. We therefore refer to these pauses as *opening-auction extensions*. A pause at the 12:30 PM open likewise extends an opening auction, although the stock has already traded in the morning session. A pause during the session, by contrast, interrupts an ongoing continuous market and replaces continuous matching with an auction. Sections 4 and 5 study opening-auction extensions at the morning open; Section 6 studies pauses that interrupt continuous trading, together with pauses at the afternoon open.

Figure 1 illustrates a pause at the open, using Toyota Motor Corporation stock on March 13, 2020, when market volatility surged due to the COVID-19 pandemic. At 9:00 AM, the matching price was 5,850 JPY, well below the lower edge of the initial price-limit range indicated by the dotted line. The opening was therefore delayed, and trading began at 9:09 AM at 5,926 JPY, after the special quote had been updated three times. The order book remained active throughout the pause, and the matching price changed as orders were entered, amended, or cancelled.

[Insert Figure 1 here]

A distinctive feature of the TSE is that the range step L is set as a fixed yen amount for each level of the reference price, rather than as a percentage of it. As a result, the percentage distance from the reference price to the triggering threshold varies discontinuously and non-monotonically across price levels. Section 4.1 uses this variation to construct our instrument.

[Insert Table 1 here]

In the United States, the LULD Plan, implemented from 2013, sets price bands of 5% or 10%, depending on the stock, around the average transaction price over the preceding five minutes. When the best quote remains at a band, the primary listing exchange declares a trading pause of at least five minutes, after which trading resumes through a reopening auction. The Plan serves a purpose similar to that of TSE pauses, but its bands are wider, and the primary listing exchange's opening auction is exempt from it. A large number of LULD pauses occur in securities outside the S&P 500 and Russell 1000 (Moise and Flaherty, 2017).

Table 2 compares single-stock pause rules across major exchanges: the TSE, NYSE and Nasdaq, and the London Stock Exchange (LSE). Unlike exchanges in the U.S. and the U.K., which use percentage-based triggers, the TSE uses fixed yen thresholds, so the sensitivity of the trigger varies with the stock's price level. A TSE pause can also end as soon as counter-orders allow the auction to clear, whereas

[2]The TSE also operates a separate mechanism, the sequential trade quote (*renzoku-yakutei-kehai*), introduced in 2010. It intervenes when a sequence of executions would move the price by more than twice the range step and switches the stock to auction matching for a nominal one minute. Our analysis concerns special quotes only.

[3]In our main sample, 37 morning-open pauses last until the close; they account for the 360-minute maximum in Table 4.

pauses on the LSE and other major European venues, commonly referred to as volatility interruptions, switch the stock to an auction that lasts for at least a set period (Zimmermann, 2014). Finally, TSE pauses can occur during the opening auction, whereas the U.S. rule does not apply to it. Unlike complete trading halts, which exchanges impose at their discretion, these mechanisms are triggered automatically by price movements and typically last only minutes. During continuous trading, they replace continuous execution with an auction; at an opening auction, a pause extends the auction. In either case, orders can still be entered, so price discovery can continue during the pause.

[Insert Table 2 here]

3 Data and Preliminary Analysis

3.1 Data and Sample

We use proprietary order book data provided by the TSE. The TSE operates as a fully electronic, order-driven market, where trades are matched based on incoming buy and sell orders, rather than through market makers as in quote-driven markets. This market structure enables direct observation of each order's direction, price, and timing, which we exploit in Section 5 to decompose the order flow during trading pauses and characterize the order-flow patterns that help explain their effects on market quality.

Our dataset captures the full lifecycle of orders, including submissions, executions, modifications (in price or quantity), and cancellations, across both pre-opening and continuous trading sessions. Each order record is time-stamped to the millisecond and contains detailed attributes such as order type, buy/sell indicator, order size, limit price, and flags indicating trading pauses. The data cover all First Section stocks, together with J-REITs, from January 14, 2014, to March 31, 2021; we refer to this set of securities as the full First Section universe.

To focus on economically important securities, our main analysis uses the 500 largest of these securities, including both common stocks and J-REITs (hereafter, the top-500 sample). We rank these securities by their market capitalization on their last trading day in this universe during the previous calendar year and keep the top 500 for the whole year. These are the largest and most actively traded securities in the Japanese market, which makes them well suited for analyzing the effects of single-stock trading pauses. As the robustness checks in Section 4.2 show, our results are not sensitive to the specific threshold used to define the sample. The resulting panel comprises approximately 873,000 stock–date observations over 1,763 trading days; roughly 9% of these observations involve a trading pause at the morning open.[4]

3.2 Frequency, Determinants, and Durations of Trading Pauses

We begin our empirical analysis with descriptive statistics. Table 3 reports when trading pauses occur within the trading day. Trading pauses are overwhelmingly concentrated at the morning open, with 77,528 newly occurring pauses, corresponding to 8.9% of all stock–date observations and 86% of all pauses in the sample. By contrast, pauses at other times are much less frequent: they occur on 0.6% of stock–dates during morning continuous trading, 0.6% during afternoon continuous trading, and only 0.3% at the afternoon open. Splitting by pause direction shows that morning-open pauses are nearly symmetric between sell-side and buy-side imbalances (4.4% vs. 4.5%), whereas buy-side pauses are more frequent during the continuous sessions and sell-side pauses at the afternoon open. Given this

[4]Three analyses require events that are rare among large stocks and therefore use the full First Section universe (about 2,100 securities on a typical day; the First Section was the exchange's main market segment until April 2022): the analysis of pauses during continuous trading in Section 6, the magnet-effect test of Section 6.4, and the spillover analysis of Section 5.4, which excludes J-REITs because the TSE's industry classification, used to define peers, does not cover them. The opening results are consistent across the two samples. In the full First Section universe, the morning-open effect is -0.470 (s.e. 0.063). Estimated separately by size group in this universe, it is -0.311 (s.e. 0.051) for the top-500 stocks, the same as the baseline estimate in Section 4, and -0.417 (s.e. 0.077) for the remaining stocks (not tabulated).

concentration, our main analysis focuses on pauses at the morning open, which extend the opening auction (Section 2).

[Insert Table 3 here]

To better understand the determinants of trading pauses at the morning open, Figure 2 reports average marginal effects from a probit model of pause incidence. The likelihood of a pause increases sharply on days following earnings announcements (or forecast revisions) released at or after the 3:00 PM close and following large absolute overnight market moves, and it is also higher when previous-day intraday volatility is elevated. In contrast, *DistanceToSQ* has a negative marginal effect, indicating that trading pauses are more likely when the reference price is closer to the triggering threshold. Formally defined in Section 4.1, *DistanceToSQ* measures the percentage distance from the reference price to the threshold that triggers a pause. We report the underlying coefficient estimates in Table A.1.

[Insert Figure 2 here]

We next examine the duration of trading pauses once they are triggered at the morning open. Figure 3 plots the distribution of pause durations, with counts on a base-10 logarithmic scale. The distribution is highly right-skewed: many pauses end within the first few minutes, while a non-trivial fraction persist for longer. A salient feature is the pronounced clustering at three-minute multiples (e.g., 3, 6, and 9 minutes). This pattern reflects the institutional design: the special quote is updated every three minutes, and each update widens the range within which the auction can clear, mechanically raising the probability that a pause ends at those times. Unlike a LULD pause, which lasts at least five minutes, a pause on the TSE can end at any time once counter-orders allow the auction to clear. The median duration is 3 minutes, and longer pauses are infrequent.

[Insert Figure 3 here]

To assess whether pause durations are predictable from information available at 9:00 AM, we examine how duration varies with the opening order imbalance, measured as the absolute difference between the 9:00 AM matching price and the reference price in multiples of the range step L. Figure 4 plots the cumulative distribution of pause durations by this multiple. The curves jump at three-minute intervals, when the special quote is updated, and shift with the imbalance: pauses last longer when the initial imbalance is larger, and the imbalance alone explains 23% of the variation in log duration. Pauses also tend to end shortly after the start of the pause or a scheduled update: 84% end within the following 30 seconds. To a first approximation, a pause ends no later than the update at which the special quote reaches the 9:00 AM matching price, although orders arriving during the pause can move the matching price. Participants can therefore anticipate, to some extent, how long a pause will last; we return to the implications of this predictability in Section 6.3.

This pattern follows from the design of the mechanism. When imbalances are small, pauses tend to be short, and the stock opens promptly; when imbalances are large, the widening range gives liquidity providers more time to submit offsetting orders before the auction clears. Whether this additional time improves the outcome of the auction is the question we examine in Section 4.

[Insert Figure 4 here]

Finally, Table 4 reports summary statistics for trading pauses at the morning open, along with key variables of interest used in our analysis. The first two rows provide descriptive statistics related specifically to trading pauses, while the remaining rows summarize the full sample. We define the daily frequency of trading pauses as the proportion of stocks experiencing a pause on a given day. This frequency varies substantially across trading days, indicating that market-wide shocks play an important role in triggering pauses. On March 13, 2020, when the COVID-19 pandemic caused a sharp market downturn, trading pauses were triggered for nearly all stocks at the morning open.

Our sample consists of relatively large stocks, with an average market capitalization of 1,036 billion JPY, approximately 9.35 billion USD at the exchange rate at the end of the sample period. Among stocks that experienced a trading pause, the average market capitalization is slightly higher, 1,074 billion JPY, indicating that trading pauses are not concentrated among the smaller firms in the sample.

[Insert Table 4 here]

3.3 Price Discovery During the Pre-opening Phase

Before turning to our instrumental-variable analysis, we examine price discovery during the pre-opening phase. Specifically, we ask how this process differs depending on whether a trading pause is triggered at the morning open and, if so, how long it lasts. To address this question, we employ the unbiasedness regression approach, following Biais et al. (1999), Barclay and Hendershott (2008), and Hautsch and Horvath (2019), and evaluate price efficiency by the slope coefficient β in the following regression:

$$\frac{v_{i,t} - E_0[v_{i,t}]}{E_0[v_{i,t}]} = \alpha + \beta \frac{P_{i,\tau} - E_0[v_{i,t}]}{E_0[v_{i,t}]} + \varepsilon_{i,t}, \quad \tau < t \tag{1}$$

where $v_{i,t}$ represents the fundamental value of stock i at time t, proxied by the closing price of the current day, and $P_{i,\tau}$ denotes the price at an earlier time τ on the same day; in this section, it is the matching price during the pre-opening phase. $E_0[v_{i,t}]$ refers to the unconditional estimate of the fundamental value, which we approximate using the base price. Both sides are therefore simple returns relative to the base price. A slope of $\beta = 1$ is the benchmark for unbiasedness in this regression, while $\beta < 1$ and $\beta > 1$ indicate overshooting and momentum, respectively.

Figure 5 illustrates the progression of price discovery during the pre-opening phase by plotting the estimates of β from repeated regressions of Equation (1), using matching prices derived from pre-opening order book data. Three patterns stand out. First, price efficiency at 9:00 AM is lower when trading pauses are triggered, and more so when pauses last longer. This pattern is consistent with incomplete information aggregation at 9:00 AM, but does not by itself establish that participants withhold orders. Second, price discovery continues during pauses and accelerates just before trading resumes, resembling the pre-opening price discovery documented for days without pauses (Biais et al., 1999; Bellia et al., 2020). Third, stocks that experience trading pauses exhibit higher price efficiency between 8:54 AM and 8:59 AM, suggesting that the information shocks that trigger pauses draw earlier attention from market participants. We revisit these patterns with an instrumental-variable approach in Section 4.3.

[Insert Figure 5 here]

4 Empirical Results Using Instrumental Variables

4.1 Instrumental Variables

The endogeneity of pause triggering confounds naive comparisons. Figure 6 illustrates this concern: stocks experiencing trading pauses exhibit higher volatility over the first 30 minutes of trading (post-open volatility), but this correlation reflects the underlying information shocks rather than a causal pause effect.

[Insert Figure 6 here]

To identify the causal effects of trading pauses, we instrument the trading-pause indicator with the percentage distance from the reference price to the triggering threshold, which we denote by *Distance-ToSQ*. Here, SQ stands for *special quote*, the quote the TSE displays during a pause (Section 2). At the

morning open, the reference price is the exchange's base price (Section 2). Let $P_{i,t}^{ref}$ denote the reference price and $L(P_{i,t}^{ref})$ the range step (in JPY) set by the TSE's schedule in Table 1; then

$$DistanceToSQ_{i,t} \equiv \frac{L(P_{i,t}^{ref})}{P_{i,t}^{ref}} \times 100. \tag{2}$$

Because the reference price is fixed before the open and the schedule is a published, yen-denominated step function, the instrument for the morning open is fully determined before any order of the trading day is submitted. This predetermination is specific to the opening auction; the instrument for pauses during continuous trading must reference a same-day transaction, and Section 6.1 and Appendix B discuss its construction.

Figure 7 visualizes how *DistanceToSQ* varies with P^{ref}. Because $L(\cdot)$ is specified in discrete JPY increments that depend on the reference-price level, *DistanceToSQ* changes discontinuously and non-monotonically as P^{ref} crosses the price-level thresholds of the schedule; in our sample, it lies between 1.5% and 2.7% for 90% of stock–dates (Table 4).

[Insert Figure 7 here]

We argue that *DistanceToSQ* plausibly satisfies the relevance and exogeneity conditions for a valid instrument. Relevance holds by construction: a smaller *DistanceToSQ* places the reference price closer to the triggering threshold, implying a higher likelihood of a trading pause. Consistent with this intuition, Figure 8 shows that observations with trading pauses are concentrated at smaller *DistanceToSQ* values than observations without pauses. Table A.2 reports the first-stage regression of the trading-pause indicator on *DistanceToSQ*. The estimates confirm a strong relationship.

[Insert Figure 8 here]

For exogeneity, we require that *DistanceToSQ* affect outcomes only through the trading-pause indicator, conditional on covariates. Table A.3 reports reduced-form regressions of post-open volatility on *DistanceToSQ*, showing that larger *DistanceToSQ* is associated with higher post-open volatility. The positive sign is what one would expect if trading pauses reduce post-open volatility: a larger *DistanceToSQ* lowers the probability of a pause, so the volatility-reducing effect is less likely to materialize. While alternative channels can never be ruled out completely, we argue that *DistanceToSQ* plausibly satisfies exogeneity for three reasons.

First, *DistanceToSQ* is a deterministic function of the predetermined reference price and the TSE's fixed range-step schedule. As a result, it is set before trading begins and cannot be influenced by real-time order flow, investor sentiment, or contemporaneous market-wide volatility.

Second, it is difficult to motivate a direct effect of the absolute price level on the post-open outcomes we study. Our results are robust when we restrict the sample to the largest stocks on the TSE and when we exclude low-priced stocks. Moreover, although other price-level-dependent microstructure features (most notably tick size) could in principle affect liquidity and volatility, we explicitly control for relative tick size, and the estimates remain stable. To the best of our knowledge, standard microstructure theories do not imply a systematic relationship between absolute price levels and the outcomes considered here once liquidity and volatility controls are included. Empirically, the log reference price enters the baseline volatility regression with a statistically significant but economically negligible coefficient (Table 5, Column (3)): a doubling of the price level is associated with about 1% higher post-open volatility, compared with the reduction of roughly 27% associated with a pause.

Third, *DistanceToSQ* varies non-monotonically and discontinuously with the reference price, generating substantial dispersion even among stocks with similar price levels. In our preferred specifications, we include both date and firm fixed effects, so identification comes from residual within-firm and within-date variation in *DistanceToSQ* beyond time- and firm-specific factors. In particular, the discontinuities

in the range-step schedule generate meaningful changes in *DistanceToSQ* when a stock's reference price crosses schedule thresholds, while leaving nearby prices otherwise similar.[5]

To further support the identification strategy, Figure A.1 plots the distributions and pairwise relationships among market capitalization, intraday realized volatility, share turnover, the reference price, and *DistanceToSQ*. The figure reveals no clear association between *DistanceToSQ* and these characteristics across stocks; within firms, *DistanceToSQ* is weakly correlated with some lagged characteristics, such as lagged volatility and market capitalization, which the baseline specification controls for.

4.2 Post-Open Volatility

Hereafter, we present and discuss our main empirical results using the instrumental variable approach. We specify the primary regression model as follows:

$$\xi_{i,t} = \alpha + \beta D_{i,t} + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}, \tag{3}$$

where $\xi_{i,t}$ denotes the outcome variable of interest, $D_{i,t}$ is an indicator variable for a trading pause at the morning open, $\mathbf{X}_{i,t}$ is a vector of control variables, and $\varepsilon_{i,t}$ is the error term. Among the covariates, we include the relative tick size, defined as the ratio of the minimum tick size to the opening price, to account for potential influences of price granularity. We also include firm and date fixed effects to control for unobserved heterogeneity across firms and trading days; the date fixed effects absorb market-wide overnight news.

The first part of our empirical analysis examines the impact of trading pauses on post-open volatility. Specifically, we analyze the log of realized volatility, computed from minute-by-minute bid–ask midpoint prices over the 30-minute window after the stock's opening trade, whether or not a pause occurs; we refer to outcomes measured over this window as post-open outcomes.[6] We follow Abad and Pascual (2010) in selecting this window, though our results are robust to alternative specifications, as we show below.

To address potential endogeneity, we estimate Equation (3) by two-stage least squares, instrumenting $D_{i,t}$ with $\mathit{DistanceToSQ}_{i,t}$ as described in Section 4.1. Unless otherwise noted, standard errors are two-way clustered by firm and date.[7] The instrument is strong: the first-stage F-statistics in Table 5 range from 606 to 672, well above the threshold of 100 beyond which conventional t-based inference is reliable (Andrews et al., 2019; Lee et al., 2022). For single-instrument estimates with weaker first stages, we add Anderson–Rubin 95% confidence sets, which remain valid when the instrument is weak.

Table 5 reports OLS and IV estimates with firm and date fixed effects. Across IV specifications, the coefficient on $D_{i,t}$ is significantly negative both with and without controls. We refer to Column (3), which includes the richest set of covariates, as our baseline specification. The corresponding estimate implies that an opening-auction extension reduces post-open volatility by approximately 27% on average $\big(= 1 - \exp(-0.311)\big)$.

In contrast, the OLS coefficient is significantly positive, consistent with the endogenous triggering discussed above. The IV estimate captures the local average treatment effect (Imbens and Angrist, 1994) for stock–dates whose pause status at the open is shifted by $\mathit{DistanceToSQ}_{i,t}$. Under monotonicity, which requires that a larger *DistanceToSQ* never makes a pause more likely, these compliers are stock–dates whose pre-open order imbalance implies a price move comparable to the range step, roughly 1.5% to 2.7% of the reference price. The estimate therefore concerns moderate overnight shocks rather than the extreme shocks that trigger a pause under any range step.

[5] In unreported analyses, we further confirm robustness using a fuzzy regression-discontinuity-style approach. Specifically, we compare observations in neighborhoods around jump points in *DistanceToSQ* with those outside such neighborhoods, adding jump-point fixed effects. The results are qualitatively and quantitatively similar across a wide range of neighborhood definitions.

[6] The window runs from the minute of the stock's opening trade through the 30th minute after it, and realized volatility is the standard deviation of the one-minute returns within it; traded value, quoted spreads, and depth in Table 4 are computed over the same interval. We exclude an observation if its window contains a further trading pause, other than the one that ends at the opening trade. We also exclude windows that overlap with the lunch break or the market close.

[7] Our inference is also robust to Driscoll–Kraay standard errors. First-stage F-statistics are cluster-robust Wald statistics, and each table note states the clustering. For the baseline estimate, the Anderson–Rubin confidence set, $[-0.41, -0.21]$, coincides with the Wald interval when both are rounded to two decimal places.

Because a pause delays the stock's opening trade, the 30-minute window of a paused stock starts later in clock time, when market-wide volatility is typically lower. When volatility is measured relative to market-wide volatility over the same minutes, the reduction is about 21%, so it remains substantial after accounting for the timing of the window.[8]

[Insert Table 5 here]

Table 6 shows that the reduction is robust to the choice of sample and of volatility measure. It holds among TOPIX 100 constituents, the largest and most liquid firms on the TSE, although it is somewhat smaller there (-0.241), and among stocks priced at 500 JPY or more, for which *DistanceToSQ* cannot take the extreme values that arise at low prices. It also holds when volatility is measured in levels, from transaction prices, by the largest price excursion or the high–low range within the window, and over the full trading day. All estimates are significant at the 1% level, except two that are significant at the 5% level.

[Insert Table 6 here]

Next, we examine the effects on post-open volatility across different time horizons. Figure 9 plots the estimated coefficient on the trading-pause indicator from the baseline IV specification as we vary the start minute used to construct the post-open volatility measure. The point estimates are uniformly negative across all start minutes. The magnitude is largest when the window begins immediately after the open (-0.31 at start minute 0) and attenuates gradually as the start minute increases (about -0.07 at 90 minutes). The shaded range denotes the corresponding 95% confidence interval, which remains below zero for windows starting up to 75 minutes after the open and includes zero thereafter. The volatility-reducing effect is thus strongest in the immediate aftermath of the open and fades over the following hour or so.

[Insert Figure 9 here]

A further check addresses the concern that overnight information confounds the opening estimates. We add to the baseline specification an indicator for corporate disclosures, such as earnings and dividend announcements, released at or after 3:00 PM on the preceding trading day. The volatility effect is essentially unchanged (-0.314, compared with -0.311; not tabulated).[9]

4.3 Price Efficiency and Post-Open Liquidity

Next, we revisit the analysis of price efficiency previously examined without an instrumental-variable approach in Figure 5. To evaluate the causal effect of trading pauses on price efficiency, we extend the unbiasedness regression in Equation (1) by introducing a trading-pause indicator and its interaction term:

$$\frac{v_{i,t} - E_0[v_{i,t}]}{E_0[v_{i,t}]} = \alpha + \beta^{SQ} D_{i,t} + (\beta + \gamma D_{i,t}) \frac{P_{i,\tau} - E_0[v_{i,t}]}{E_0[v_{i,t}]} + \varepsilon_{i,t}, \tag{4}$$

where $v_{i,t}$ and $E_0[v_{i,t}]$ are defined as in Equation (1). We estimate the model using prices at two times: 9:00 AM and the time of the stock's opening call auction. In the 9:00 AM specification, $P_{i,\tau}$ is the opening price for non-paused stocks, which open at 9:00 AM, and the matching price for paused stocks. Our primary parameter of interest is γ, which captures how trading pauses affect price efficiency.

[8] The outcome is the log ratio of the stock's realized volatility to market-wide realized volatility over the same minutes, computed as the equal-weighted average, across top-500 stocks, of squared one-minute midpoint returns; the IV estimate is -0.230 (s.e. 0.051; not tabulated).

[9] We also examine whether paused and non-paused stocks differ in their post-open response to overnight market news. In a regression of the post-open 30-minute return, the interaction of the pause indicator with the overnight market return, with both terms instrumented by *DistanceToSQ* and its interaction, is small and insignificant (-0.029, s.e. 0.039): we detect no statistically significant difference between paused and non-paused stocks in how they track market-wide overnight news after the open. This test does not separately identify learning from market movements during the extension.

Panel A of Table 7 reports estimates of Equation (4) using both OLS and IV specifications.[10] Following the existing literature on the unbiasedness regression, the baseline specifications include no firm or date fixed effects; Table A.4 in Appendix A shows similar estimates with date fixed effects and two-way clustered standard errors. The OLS estimates of β are 0.961 in both specifications, indicating modest intraday reversal, consistent with Lou et al. (2019); the corresponding IV estimates are 1.044 and 0.985.

The OLS and IV estimates differ notably in the 9:00 AM specification (Column (1)). The OLS estimate of γ is -0.073, whereas the IV estimate is larger in magnitude at -0.202. This gap is consistent with attenuation in OLS: trading pauses are triggered by large information shocks that also move prices toward fundamentals, leading OLS to understate the lower price efficiency of the 9:00 AM matching price. The gap may also reflect the populations to which the two estimates apply: the IV estimate concerns compliers facing moderate overnight shocks (Section 4.2), whereas OLS averages over all pauses. In the opening-price specification (Column (2)), the OLS estimate of γ is positive and significant (0.020), but the IV estimate is -0.020 and statistically insignificant. Taken together, the IV results imply that the 9:00 AM matching price of paused stocks is significantly less efficient, while any difference at the actual open is not detectable.

This pattern is consistent with the view of the pause as an auction extension. The 9:00 AM matching price aggregates the orders present at 9:00 AM but not the offsetting liquidity that arrives during the extension. Consistent with this reading, Figure 5 shows that price discovery continues during pauses and that, once trading starts, price efficiency becomes comparable regardless of whether a trading pause was triggered.

[Insert Table 7 here]

Panel B of Table 7 makes the same comparison directly for paused stock–dates, for which both the 9:00 AM matching price and the eventual opening price are observed. It reports the unbiasedness slope of the close-to-close return on (i) the close-to-9:00 AM return, computed with the matching price, and (ii) the close-to-open return, computed with the actual opening price. The slope is 0.888 for the matching price, so that about 11% of the 9:00 AM price movement is subsequently reversed by the close, and 0.981 for the opening price, closer to one although still statistically below one. In a regression that includes both returns, the coefficient on the close-to-open return is 0.991 and that on the close-to-9:00 AM return is -0.010: once the opening price is known, the 9:00 AM matching price adds no detectable predictive information about the closing price. Because we do not observe how the prices of paused stocks would have evolved without an extension, the Panel B comparison is descriptive within the paused sample rather than causal; we read it as price-level evidence consistent with the volatility results.

Next, we examine the effects of trading pauses on market liquidity, using the same empirical framework as in the post-open volatility analysis. Specifically, in Equation (3) we replace the dependent variable with the average quoted bid–ask spread, computed from minute-by-minute order book data from the start of trading through 30 minutes afterward, so that a lower value indicates higher liquidity.

Table 8 shows that both OLS and IV estimates imply a statistically significant decline in quoted spreads following trading pauses. The magnitudes, however, differ sharply: in the baseline specification (Column (3)), the OLS estimate is -0.356 basis points, whereas the IV estimate is -6.249 basis points. Unlike in the volatility regressions, where the OLS and IV estimates have opposite signs, the OLS estimate here has the same sign as the IV estimate but is much closer to zero. This gap is consistent with information shocks both making pauses more likely and widening spreads, which would attenuate the negative OLS association; the larger IV estimate is consistent with the decline in volatility.[11]

[10] The TSE does not allow market participants to submit orders beyond the specified daily price limits. Consequently, matching prices cannot be observed outside these bounds. To mitigate potential bias from missing observations at the extremes, we exclude observations for which the absolute close-to-9:00 AM return exceeds the minimum percentage distance to the daily price limits observed in the sample universe.

[11] Because quoted spreads are bounded below by the tick, the spread in basis points is approximately proportional to the relative tick size, which Table 8 therefore controls for linearly. With the log spread and the log relative tick size instead, the IV estimate is -0.23 ($t = -2.2$), a reduction of roughly 20%.

Table A.5 in Appendix A further reports IV estimates for traded value, tick-adjusted depth, and the value executed in the opening call. While the first two measures increase significantly in specifications without additional controls, the effects become statistically insignificant once market capitalization, share turnover, and lagged volatility are included. Taken together, the evidence of improved liquidity is clearest in quoted spreads; the estimates for trading activity and depth do not establish a corresponding change in the market's capacity to absorb volume. We return to this point when decomposing the volatility reduction in Section 5.1. We further investigate the mechanisms of liquidity provision in Section 5.2 using detailed order flow data.

[Insert Table 8 here]

5 Mechanisms, Heterogeneity, and Spillovers

5.1 The Source of the Volatility Reduction

The volatility reduction is consistent with reduced transitory price variation: post-open path variation falls more than the squared net price change over the same window, and we detect neither a statistically significant decline in trading activity nor evidence that volatility is displaced to later periods. A lower 30-minute realized volatility after a pause could arise from several mechanisms with different welfare implications: the auction may clear at a price that requires less subsequent correction, trading may simply be suppressed for a while, quote updating may become sparser so that midpoint volatility is mechanically lower, or the same volatility may be displaced to later in the day. We separate these explanations by re-estimating the baseline IV specification with alternative dependent variables.

We start from an identity that links the variation of the price path to its net movement. Let r_m denote the one-minute midpoint log returns in the 30-minute window and $r_{\text{tot}} = \sum_m r_m$ the net return over the window. The path variation is $\sum_m r_m^2$, and the variance ratio $\text{VR} = r_{\text{tot}}^2 / \sum_m r_m^2$ measures how much of the path's variation is net movement rather than round trips: for a random walk it is one on average, and it falls as prices zigzag around their eventual level. We estimate the effect of the pause on $\ln \sum_m r_m^2$ (the path variation), $\ln r_{\text{tot}}^2$ (the net movement), and $\ln \text{VR}$, which is their difference (Table 9).

[Insert Table 9 here]

Path variation falls more than the net movement. On the common sample, the log path variation falls by 0.575 ($t = -5.7$), whereas the decline in the log squared net return, 0.212, is not statistically significant ($t = -1.6$), so the log variance ratio rises; the outcome-specific samples give the same pattern. This pattern is consistent with reduced transitory price variation, that is, less back-and-forth movement of prices around their eventual level; the decomposition does not, however, separately identify permanent and transitory components. The maximum price excursion from the opening price within the window also falls (Table 6, Column (5)).

No statistically detectable decline in trading activity. If volatility fell because trading stopped, traded value would fall with it. Under the baseline specification, traded value, tick-adjusted depth, and the value executed in the opening call do not change significantly (Table A.5), although these estimates do not rule out economically meaningful reductions in activity. Together with the narrower spreads of Table 8, the evidence is more consistent with better price alignment than with a greater capacity to absorb trading volume.

Measurement. Because midpoint volatility depends on quote updating, a pause that made quotes sparser could lower it mechanically. Transaction-price realized volatility, which does not depend on quotes, also falls (-0.253, $t = -4.9$; Table 6, Column (4)). If quotes merely went stale while trading continued, the ratio of transaction-price to midpoint volatility would rise; it shows no increase that is significant at the 5% level. These results make an explanation based solely on stale quotes unlikely. However, the share of return intervals without a midpoint change rises after a pause, mainly among

coarse-tick stocks, so part of the measured decline in midpoint volatility could reflect sparser quote updating; we regard this as the main measurement limitation of the analysis.[12]

No evidence of displacement. Figure 9 shows that the reduction decays as the measurement window is shifted later into the session but does not reverse, and realized volatility over the full trading day is lower as well (Table 6, Columns (7) and (8)). Realized volatility over the full next trading day is also lower (-0.14, $t = -2.8$; not tabulated). We therefore find no evidence that the pause merely postpones volatility, as critics of trading pauses have argued it might (Kyle, 1988; Fama, 1989).

Taken together, these results are consistent with the additional auction time allowing offsetting liquidity to arrive and the clearing price to settle, which the order-flow evidence of Section 5.2 examines directly. They weigh against the alternatives, although they do not rule out some contribution from each: trading suppression predicts lower traded value, a stale-quote artifact predicts a rising ratio of transaction-price to midpoint volatility, and displacement predicts higher volatility later.

5.2 Liquidity Provision During Trading Pauses

The order-level data show that, during opening-auction extensions, counter-orders against the direction of the initial imbalance arrive throughout the pause, are mostly priced more aggressively than the current matching price, and, within complete update cycles, are mostly placed beyond the current special quote. We interpret this as liquidity provision: the extension gives market participants time to absorb information and submit orders that offset the initial imbalance. We present the evidence in order, from aggregate price patterns to order-level decompositions.

Panel A of Figure 10 plots the close-to-open return (using the actual opening price) against the close-to-9:00 AM return (using the matching price at 9:00 AM) for all trading-pause observations. If prices fully reverted during the pause, the slope would be zero; if no reversal occurred, the slope would be one. The slope of the regression through the origin is 0.859, implying a fitted reversal of about 14% of the 9:00 AM price movement by the time trading resumes. Panels B and C of Figure 10 split the observations by the sign of the close-to-open return: the slope is 0.877 for negative returns (Panel B) and 0.844 for positive returns (Panel C). Reversals are thus weaker after price declines, possibly reflecting greater uncertainty and risk aversion among liquidity providers; we return to this asymmetry in Section 5.3. The figure also shows a pronounced cluster of close-to-9:00 AM returns near -10%, which may reflect both a round-number focal point for pre-opening orders and the short-selling price restriction triggered at that decline.[13]

[Insert Figure 10 here]

Figure 11 summarizes order flow in each minute of the pause. Panel A reports, for each minute, the share of pauses in which net order volume against the pause direction exceeds that along it, with cancellations counted as negative volume. For limit orders, this share exceeds one-half in every minute, consistent with liquidity provision that offsets the initial imbalance; for market orders, it stays below one-half throughout, consistent with momentum-driven demand. Panel B classifies new orders as active if they are market orders or limit orders priced more aggressively than the current matching price, and as passive otherwise. Against-direction orders account for more than half of classifiable new-order volume in every minute (51–58%, averaged across pauses with equal weight), and most of them are active, whereas along-direction orders show a more balanced mix. Liquidity providers thus do not merely queue behind the matching price; they price their orders aggressively enough to shift it toward reversal. Cancellations of along-direction orders also account for a substantial part of the reversal pressure (Figure A.2 in Appendix A).

[12] The IV estimate for the share is 0.39 ($t = 5.6$; not tabulated); because the share is bounded, the estimate is informative mainly about direction. Estimated separately, it is 0.35 ($t = 8.0$) for stocks with a relative tick size above the sample median and -0.05 ($t = -1.6$) below it. Each subsample estimate has its own first stage and compliers, so the full-sample estimate need not lie between them.

[13] Since November 2013, and thus throughout our sample period, once a stock's price falls by 10% or more from its base price, short sales are subject to a price restriction similar to an uptick rule for the rest of that day and the following trading day. We do not separate the two explanations.

[Insert Figure 11 here]

Table 10 relates these orders to the three-minute widening of the price-limit range, using complete update cycles within the first 30 minutes of morning-open pauses. One might expect liquidity providers to time their submissions to the updates, or to hold back until the expected resumption and then submit at the newly updated special quote. Nothing in the mechanism requires either, because the special quote is not a limit on order prices (Section 2). Two facts stand out. First, against-direction limit orders arrive throughout the cycle, with only modest clustering at its two ends: 18.8% arrive within 30 seconds after an update and 18.5% within 30 seconds before the next one, compared with 16.7% under uniform arrival (Panel A).

Second, against-direction limit orders are placed ahead of the special quote rather than at it (Panel B). In the first cycle, 78% are priced beyond the special quote by more than half a range step and only 15% near it, and about half are priced beyond the level the special quote will reach at the next update. The share priced beyond the special quote rises with the size of the initial imbalance, from 64% for the smallest imbalances to 85% for the largest (not tabulated), consistent with against-direction orders clustering around the 9:00 AM matching price. The prices of these orders typically do not change from one update to the next, whereas the special quote moves by one range step at each update (not tabulated). The mechanically updated special quote thus moves toward orders that are already in place.

Arrival immediately before the opening execution is a separate question. New against-direction limit orders are concentrated in the final seconds before the opening execution: 3.7% of these orders and 5.4% of their volume arrive in the final 10 seconds, compared with 2.3% and 2.2% if arrivals were uniform over the pause, and the concentration is weaker for along-direction orders (not tabulated). Most counter-orders nevertheless arrive earlier.[14] This concentration does not identify strategic withholding: an incoming counter-order can itself allow the auction to clear, and we do not observe when these orders would have arrived without the extension. Moreover, the order-level records used in this analysis begin with the posting of the special quote, so they do not reveal orders withheld before 9:00 AM. Section 6.3 discusses what the predictability of pause durations implies for our IV estimates.

Taken together, the order-level evidence is consistent with a simple account of how opening-auction extensions work: the extension gives liquidity providers time to submit aggressively priced counter-orders, placed ahead of the widening range, that help absorb the initial imbalance and move the matching price toward reversal. These patterns are descriptive, and they do not separately identify learning about market-wide information from the arrival of offsetting liquidity, but they are consistent with the reversals in Figure 10 and with the decline in volatility and the narrowing of spreads documented in Section 4.

These results also bear on the design of the mechanism, in which the update schedule is a key parameter. The schedule shapes when the auction clears: durations cluster at three-minute intervals (Section 3), and counter-orders are placed ahead of the widening range. The evidence does not, however, establish whether the three-minute schedule clears the auction at the optimal time: because our instrument varies pause incidence under a fixed schedule, it cannot compare alternative update intervals or range steps.

5.3 Heterogeneity in the Effectiveness of Trading Pauses

The volatility reduction is larger for stocks with lower share turnover and smaller after the stock's overnight price declines, after large market-wide overnight moves, and for stocks that have recently been volatile or falling; after earnings announcements, its point estimate is close to zero. We examine two dimensions of liquidity provision that can shape the effect of additional auction time:

Hypothesis 1: The volatility reduction is larger where baseline trading activity is low, because offsetting orders that arrive during the extension have a larger impact relative to the orders already present.

[14] The sample consists of pauses with at least one update that end within 30 minutes and whose order records extend to within one second of the opening execution. For along-direction orders, the corresponding shares are 2.8% and 3.4%, compared with 2.2% for both under uniform arrival.

Hypothesis 2: The volatility reduction is smaller where information or inventory risk makes liquidity providers reluctant to supply offsetting orders.

To this end, we interact the trading-pause indicator $D_{i,t}$ with a binary variable $Z_{i,t}$ that captures each dimension of interest:

$$\xi_{i,t} = \alpha + (\beta + \delta Z_{i,t})D_{i,t} + \zeta Z_{i,t} + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}. \tag{5}$$

We instrument $D_{i,t}$ and $Z_{i,t}D_{i,t}$ with $DistanceToSQ_{i,t}$ and its interaction with $Z_{i,t}$. Tables 11–13 report IV estimates with post-open volatility as the dependent variable; the corresponding OLS estimates and IV estimates for quoted spreads are in Tables A.6–A.11. The indicator variables are defined in the table notes.

Table 11 uses indicators of size and liquidity. We detect no statistically significant difference in the volatility reduction by TOPIX 100 membership or by market capitalization below the daily median, so pauses reduce volatility for large and liquid stocks as well. The spread response differs, however: in Table A.7, the narrowing of spreads is not observed among TOPIX 100 stocks (an implied point estimate of about $+5$ basis points) and is larger for stocks with market capitalization below the daily median (about -11 basis points, compared with -2 for the others). The interaction with the Low Turnover indicator, which equals one if the stock's previous-day share turnover is below the daily median, is significantly negative: the reduction is larger for stocks with lower turnover, consistent with Hypothesis 1.

[Insert Table 11 here]

Table 12 uses indicators of overnight return characteristics. The interaction with the Negative Overnight Return indicator is significantly positive: pauses reduce volatility less after the stock's price has fallen overnight, and this asymmetry remains when market-wide declines are controlled for (Column (2)). The weaker effect after price declines matches the weaker reversals after declines in Figure 10; liquidity providers may be more hesitant to lean against a stock whose price is falling. The interaction with the Large Overnight Market Return indicator is also significantly positive, consistent with Hypothesis 2: market-wide shocks may limit the ability of liquidity providers to stabilize individual stocks. Likewise, the interaction with the Earnings Announcement indicator nearly offsets the main effect, so that the implied point estimate on the trading day following an announcement is close to zero. This pattern is consistent with the extension helping most when the imbalance reflects a temporary shortfall of offsetting orders rather than firm-specific news that the market must price.

[Insert Table 12 here]

Table 13 uses indicators of lagged volatility and recent returns, excluding the overnight return that triggers the pause. Pauses reduce volatility less for stocks with high underlying volatility and after recent declines; when the indicators are combined, the attenuation is concentrated after severe previous-day declines and declines over the previous three days.

[Insert Table 13 here]

In summary, the opening-auction volatility reduction is also present among large and highly liquid stocks, and its variation across stocks and market conditions is consistent with both hypotheses: additional auction time matters more where baseline activity is low, and liquidity providers appear more reluctant to lean against the imbalance when the stock is risky or falling, or when market-wide or firm-specific news must be priced.

5.4 Spillovers to Non-paused Stocks

The IV estimates do not detect statistically significant adverse spillovers from paused stocks to other stocks in the same industry, but their confidence sets allow economically meaningful increases in volatility. The concern is that delaying execution in some stocks at the open could impair the pricing of common factors and thereby harm *other* stocks in the same industry. We examine this concern on *non-paused*

stock–dates, asking whether exposure to paused industry peers raises opening volatility or distorts price discovery.

We measure exposure by the share of same-industry peers paused at the morning open and instrument it with the same institutional variation used throughout: the position of *peers'* base prices on the range-step schedule, summarized by the industry median of peers' *DistanceToSQ* and by the share of peers whose *DistanceToSQ* falls in the bottom decile. The sample consists of First Section stocks that are not paused at the open. Because exposure and the instruments are largely common within an industry on a given day, we cluster standard errors by industry and date. Table 14 reports the estimates, with and without controls for peers' characteristics, scaled to a 10-percentage-point increase in exposure, about one standard deviation.[15]

[Insert Table 14 here]

OLS shows a *positive* association between peer-pause exposure and own volatility, as industry-wide shocks move both. The IV estimates do not show this positive association, but their sign is not stable. The median-distance estimate changes sign once peer characteristics are controlled for, whereas the bottom-decile-share estimate is negative in both specifications. No estimate is statistically distinguishable from zero under the Anderson–Rubin confidence sets.[16] The estimates also rest on exclusion restrictions that could fail if peers' price composition, or rules that share the same thresholds, such as the daily price limits, affect own-stock outcomes directly.

No IV specification detects a statistically significant adverse spillover: every IV point estimate is negative or statistically indistinguishable from zero. The 95% Anderson–Rubin upper bounds correspond to volatility increases of about 5–7% per 10-percentage-point increase in exposure for the bottom-decile-share instrument, but of up to about 30% for the median-distance instrument with peer controls, and the point estimates change sign across specifications. The available instruments are therefore too weak to rule adverse spillovers out, and the evidence does not establish a calming effect either.

Two supplementary results are consistent with this reading. An indicator for at least one paused peer gives negative but imprecise estimates; it is nearly saturated (85% of stock–dates have a paused peer) and weakly identified. Descriptively, the unbiasedness slope of opening prices for non-paused stocks is similar for stock–dates with no paused peer and with more than 10% of peers paused (0.772 and 0.770).

A final check addresses selection into the non-paused sample. Restricting the sample to non-paused stocks conditions on each stock's own pause status, which depends on that stock's *DistanceToSQ* and may be correlated with peers' instruments when industry peers have similar price levels. When own and peer pauses are instrumented jointly among all First Section stocks, the own-pause estimate remains negative (-0.474, s.e. 0.063), and the peer-exposure estimate is close to zero (-0.004 per 10 percentage points, s.e. 0.034).

6 Pauses During Continuous Trading: A Contrast

The analysis so far concerns pauses that extend the opening auction. A minority of TSE pauses instead interrupt the continuous session, and these are the closest analogue in our setting to the pauses studied in the U.S. and European literature. This section applies the same identification strategy within the same market to ask whether the volatility reduction documented at the open carries over to pauses that suspend an ongoing continuous market. We detect no robust volatility reduction for these pauses: the estimate for price-driven pauses is close to zero with a wide confidence set, and the evidence of a reduction for order-book-evaporation pauses is marginal.

[15] J-REITs are excluded because the TSE's industry classification does not cover them; peer exposure is computed over all same-industry First Section stocks, including paused ones. The peer controls are leave-one-out industry–date means of peers' characteristics, including a flexible function of their price levels, and Table 14 also reports standard errors clustered by firm and date.

[16] A cluster-robust overidentification test of the two instruments does not reject ($J = 1.53$, $p = 0.22$), although with this little identifying variation the test has limited power. Instruments based on peers just below versus just above a threshold are too weak to be informative.

6.1 Event Types, Sample, and Design

Pauses outside the morning open are not homogeneous. Using the exchange's records of trading pauses, we distinguish three types.

- *Price-driven* pauses: an incoming order would execute at a price outside the price-limit range, and a matching price exists when the special quote is first displayed; the pause interrupts an active market.
- *Order-book-evaporation* pauses: no matching price exists when the special quote is first displayed, as when the opposing side of the book has emptied.
- *Afternoon-open* pauses (12:30 PM): the pause extends the afternoon opening auction. Institutionally, these belong with the auction extensions studied in Sections 4 and 5; we report them here because their estimation requires the intraday design of this section, and they provide a second, less precise estimate for auction extensions.

Nearly all pauses, regardless of type, end with a call auction: a call-auction execution occurs within three minutes of the end of the pause in more than 98% of the events for which auction executions are recorded separately.

The estimation sample comprises 3,438 continuous-session pause events (1,406 price-driven and 2,032 order-book-evaporation events) and 3,686 afternoon-open pauses over 2014–2021, on stock–dates with a single pause; Appendix C describes the sample filters. The outcome is log realized volatility over the 30 minutes after the pause ends.

This analysis uses the full First Section universe rather than the top-500 sample, because pauses during continuous trading occur mainly in smaller stocks: the top-500 sample contains only 56 continuous-session events. For the morning open, where both samples are large, the estimates are similar across size groups (footnote 4). Requiring a same-day reference price also excludes events that occur before the stock's first trade of the day, which limits the population that our estimates cover (Appendix C).

The design adapts the opening-auction IV to the intraday setting in two ways. First, each treated event is compared with the same stock at the same clock time on the ten preceding trading days, with firm and date fixed effects, and each treated event and its control windows receive equal total weight. In this section, "post-event" refers to the 30-minute window beginning at the end of a treated pause or at the same clock time on its matched control days. Because control windows are drawn from days without a pause, identification relies on the assumption that, within the matched sets and conditional on the controls, the instrument is unrelated to unobserved determinants of post-event volatility.

Second, for continuous-session pauses, the reference price that enters the instrument must be observed before the trigger; we reconstruct it from the last transaction that does not belong to the triggering order, at tick precision for most events (Appendix B).

The specifications use the controls of Table 5 together with pre-event realized volatility, which absorbs differences in conditions between treated and control windows. First stages are weaker than at the morning open, with first-stage F-statistics of 20 to 33, so we report Anderson–Rubin (AR) confidence sets alongside the standard errors.

6.2 Post-Event Volatility

[Insert Table 15 here]

Table 15 reports the estimates for the three types, with the morning-open estimate of Table 5 alongside for comparison. The estimate for price-driven pauses is close to zero, and its confidence set is wide in both directions: its upper bound corresponds to a volatility increase of about 62%, so moderate deterioration cannot be excluded. The estimate for order-book-evaporation pauses is negative, and its AR set barely excludes zero. However, the conventional t-test does not reject, and the AR set includes zero in two of the four robustness checks of Table A.12, so we regard the evidence of a reduction for these pauses as marginal. Both confidence sets include the morning-open estimate, so these results do

not establish a difference in causal effects across sessions, and the two continuous-session estimates are not statistically distinguishable from each other. OLS estimates on the same samples are positive (not tabulated), whereas the IV estimates are close to zero and negative, consistent with selection contributing to the positive OLS association.

The afternoon-open estimate is -0.500, significant at the 10% level by the conventional t-test and with an AR set that excludes zero, in line with the morning open. Across the robustness checks of Table A.12, the afternoon-open estimate remains negative, the price-driven estimate remains insignificant, and the AR set for order-book-evaporation pauses excludes zero in two of the four checks. Table A.13 traces the construction of the estimation samples, which Appendix C describes.

These results bear on comparisons with the U.S. market. Because LULD pauses interrupt continuous trading and the primary listing exchange's opening auction is exempt (Section 2), our continuous-session estimates address a closer institutional counterpart to LULD than our opening-auction estimates do. The opening-auction estimates should not be read as evidence on LULD pauses. Differences in event populations, institutional rules, and empirical methods further limit comparisons across studies.

6.3 Order Flow and the Predictability of Pause Duration

During price-driven pauses, against-direction orders account for only about half of classifiable new-order volume, compared with 56–57% at the opening auctions, and they are less aggressive. During order-book-evaporation pauses, along-direction orders predominate among classifiable new orders, unlike at the opening auctions.

Figure 12 compares the pause types with the morning open. The upper panel plots the minute-by-minute dominance of against-direction limit-order flow, and the lower panel the composition of classifiable new orders by direction and aggressiveness. The afternoon open reproduces the morning-open profile, with active against-direction orders the largest component. In price-driven pauses, against-direction submissions are mostly passive. In order-book-evaporation pauses, passive along-direction submissions form the largest category.

[Insert Figure 12 here]

The passive placement in price-driven pauses is consistent with counter-orders resting near the last transaction price, which provides a reference during continuous trading, rather than crossing a matching price that is still forming. These descriptive differences may help interpret the volatility estimates, but they are not causal evidence: order flow during a pause is itself an outcome of the pause, so it does not reveal what orders would have arrived without one.

Returning to the morning open, we address another concern: because pause durations are partly predictable (Section 3), market participants know roughly how much time they have. They can anticipate the update times and the latest update at which the auction is likely to clear, but not whether it will clear earlier. We view this predictability as part of the mechanism we study. Concentrating submissions toward the clearing time is how call auctions work, as in the ordinary pre-opening period (Biais et al., 1999), and our IV estimates capture the effect of extending the auction, including any adjustment in submission timing. The order-level evidence of Section 5.2 fits this view: counter-orders arrive throughout the update cycles, and their concentration just before the opening execution cannot distinguish strategic withholding from orders that themselves allow the auction to clear. Predictability may contribute to liquidity provision at the open, but we do not separately identify its role.

Whether this timing pattern carries over to pauses during continuous trading cannot be established. For pauses outside the morning open that last at least six minutes, the arrival-timing and placement statistics of Table 10 are broadly similar to those at the morning open, but only 44 of these 398 pauses are price-driven; the remainder extend the afternoon opening auction.[17]

[17] Under the definitions of Table 10, 19.6% of against-direction submissions in complete cycles after the first arrive within 30 seconds after an update and 16.1% within 30 seconds before the next one, and 71% of new against-direction orders in the first cycle are priced more than half a range step beyond the special quote. Order-book-evaporation pauses are not included in these statistics.

6.4 Magnet Effects

In our baseline tests, we find no evidence consistent with a magnet effect for pauses during continuous trading. A prominent concern with circuit breakers is that, as the price approaches a trigger, traders may accelerate execution to avoid being caught by the halt, pushing the price toward the trigger and raising volatility beforehand (Subrahmanyam, 1994; Cho et al., 2003; Chen et al., 2024). At the morning open, this channel cannot be tested, because no continuous trading precedes the pause. Pauses during continuous trading do have a pre-pause window, but comparing the minutes before a pause with control windows cannot separate anticipation from selection: conditioning on realized pauses selects unusual price and liquidity conditions, so the descriptive pre-trigger patterns in Appendix D.1 do not identify anticipation. We therefore test whether trading responds to variation in the likelihood of a pause, without conditioning on a subsequent pause.

The range-step schedule provides this variation. Because $L(p)$ is a step function of the price level, $z = L(p)/p \times 100$—*DistanceToSQ* of Equation (2) evaluated at the current price rather than the reference price—jumps at the schedule thresholds, so that, controlling for log price, otherwise similar price states face different probabilities of triggering a pause. We take continuous-session stock–minutes that follow a price move of at least 0.8% over five minutes, a threshold that does not depend on z, and regress trading behavior over the next two minutes on z, controlling for the price level, the size of the preceding move, and preceding activity (Table 16). If traders anticipate pauses, a narrower price-limit range should induce them to trade sooner and push the price toward the trigger.

Table 16 reports the results. A one-percentage-point decrease in z is associated with a 16.4-basis-point increase in the probability of a pause within five minutes, compared with a base rate of 28.8 basis points. We detect, however, no statistically significant acceleration of trading when a pause becomes more likely: the coefficient on traded value is indistinguishable from zero, price moves are, if anything, less likely to continue toward the trigger, and the response of aggressor flow toward the move is not statistically significant. Because z does not predict pause incidence among top-500 stocks (Panel B), the test is informative mainly about smaller stocks. Appendix D.2 reports robustness to move thresholds, price trims, and size groups (Table A.14); among the samples in which z predicts pause incidence, only the traded-value response at the 1.5% move threshold is significant with the sign that a magnet effect implies. A pre-trigger run-up in realized pauses (Appendix D.1) is compatible with these results, because that run-up conditions on the pause having occurred. The test concerns price-limit-range triggers tied to recent transaction prices, and is silent on static daily price limits, for which magnet effects have been documented elsewhere (Cho et al., 2003).

7 Conclusion

This paper estimates the causal effect of single-stock trading pauses on market quality, using the TSE's yen-denominated range-step schedule, which makes the percentage distance to the trigger vary across stocks with their price levels. Because most TSE pauses occur at the morning open, our main estimates concern opening-auction extensions. These extensions reduce realized volatility over the first 30 minutes of trading by roughly a quarter, consistent with reduced transitory price variation, and quoted spreads narrow without a statistically significant decline in traded value or tick-adjusted depth. Order flows suggest that the extension allows counter-orders to offset the initial imbalance; within complete update cycles, most new against-direction limit orders are priced beyond the current edge of the price-limit range. The reduction is larger for stocks with lower turnover and smaller after the stock's overnight price declines, after large market-wide overnight moves, and for stocks that have recently been volatile or falling; after earnings announcements, its point estimate is close to zero. These patterns are consistent with the effect depending on how readily offsetting liquidity can be supplied during the pause.

The evidence of a volatility reduction comes from pauses that extend a call auction. For pauses that interrupt continuous trading, we detect no robust volatility reduction: the estimate for price-driven pauses is close to zero, with a confidence set wide enough to include both sizable increases and reductions

comparable to those at the open, and the evidence of a reduction for order-book-evaporation pauses is marginal. We also detect no statistically significant adverse spillovers to non-paused industry peers, but the confidence sets allow economically meaningful increases in volatility. The baseline tests during continuous trading yield no evidence consistent with a magnet effect.

Our design identifies these effects within a regime in which pauses are already in place, without drawing the counterfactual from a period or market without pauses. In the baseline specifications, OLS associates pauses with higher subsequent volatility; IV estimates are negative for auction extensions and not robustly different from zero during continuous trading. Our own-pause IV estimates are local to observations whose pause status responds to the triggering distance. The opening analysis uses the top-500 sample; the continuous-session analysis uses selected single-pause events after a same-day transaction in the full First Section universe, predominantly in smaller stocks. Effects during extreme shocks are not separately established.

For market design, the evidence supports opening-auction extensions under the TSE's rules as a means of reducing post-open volatility, but the opening-auction results should not be extrapolated to U.S. LULD pauses. More precise estimates for continuous-session pauses could show whether their effects differ from those of opening-auction extensions. Variation in update intervals or range steps, which our design holds fixed, would make it possible to evaluate alternative designs. Trader identifiers linked to trading strategies could reveal who supplies liquidity and clarify the role of high-frequency traders (Chakrabarty et al., 2017).

Data availability

The data used in this study are proprietary and were provided by the Japan Exchange Group under a confidentiality agreement. Due to the terms of the agreement, the data cannot be shared publicly. The authors will assist with any inquiries about the data used in this study.

CRediT authorship contribution statement

Akitada Kasahara: Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Software, Visualization, Writing – original draft, Writing – review & editing. **Masahiro Yamada:** Conceptualization, Data curation, Funding acquisition, Investigation, Validation, Writing – review & editing.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used ChatGPT (OpenAI) and Claude (Anthropic) in order to improve language and readability, assist with code development for data analysis, and support manuscript formatting. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

D. Abad and R. Pascual. Switching to a temporary call auction in times of high uncertainty. *Journal of Financial Research*, 33(1):45–75, 2010.

D. Abad and R. Pascual. Holding back volatility: Circuit breakers, price limits, and trading halts. In *Market Microstructure in Emerging and Developed Markets*, pages 303–324. John Wiley & Sons Inc., 2013.

David Abad and Roberto Pascual. On the magnet effect of price limits. *European Financial Management*, 13(5):833–852, 2007.

I. Andrews, J. H. Stock, and L. Sun. Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11:727–753, 2019.

M. J. Barclay and T. Hendershott. A comparison of trading and non-trading mechanisms. *Journal of Empirical Finance*, 15(5):839–849, 2008.

M. Bellia, L. Pelizzon, M. G. Subrahmanyam, J. Uno, and D. Yuferova. Low-latency trading and price discovery: Evidence from the Tokyo Stock Exchange in the pre-opening and opening periods. Unpublished Working Paper, 2020.

B. Biais, P. Hillion, and C. Spatt. Price discovery and learning during the preopening period in the Paris Bourse. *Journal of Political Economy*, 107(6):1218–1248, 1999.

B. Chakrabarty, T. Hendershott, S. Nawn, and R. Pascual. Order exposure in high frequency markets. Unpublished Working Paper, 2017.

H. Chen, A. Petukhov, J. Wang, and H. Xing. The dark side of circuit breakers. *Journal of Finance*, 79 (2):1405–1455, 2024.

D. D. Cho, J. Russell, G. C. Tiao, and R. Tsay. The magnet effect of price limits: Evidence from high-frequency data on Taiwan Stock Exchange. *Journal of Empirical Finance*, 10(1–2):133–168, 2003.

S. A. Corwin and M. L. Lipson. Order flow and liquidity around NYSE trading halts. *Journal of Finance*, 55(4):1771–1801, 2000.

S. S. Deb, P. S. Kalev, and V. B. Marisetty. Price limits and volatility. *Pacific-Basin Finance Journal*, 45: 142–156, 2017.

E. F. Fama. Perspectives on October 1987, or, what did we learn from the crash? In R. W. Jr. Kamphuis, R. C. Kormendi, and J. W. Henry Watson, editors, *Black Monday and the Future of Financial Markets*, pages 71–82. Center for Research in Security Prices, Graduate School of Business, University of Chicago, 1989.

B. C. Greenwald and J. Stein. The task force report: The reasoning behind the recommendations. *Journal of Economic Perspectives*, 2(3):3–23, 1988.

B. C. Greenwald and J. C. Stein. Transactional risk, market crashes, and the role of circuit breakers. *Journal of Business*, 64:443–462, 1991.

Y. Hamao and J. Hasbrouck. Securities trading in the absence of dealers: Trades and quotes on the Tokyo Stock Exchange. *Review of Financial Studies*, 8(3):849–878, 1995.

L. Harris. Circuit breakers and program trading limits: What have we learned? Working paper, Wharton School Center for Financial Institutions, University of Pennsylvania, 1997.

N. Hautsch and A. Horvath. How effective are trading pauses? *Journal of Financial Economics*, 131(2): 378–403, 2019.

Guido W. Imbens and Joshua D. Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.

P. Jain, S. McFarland, and T. H. McInish. Short duration, dynamic price limits: The special quote and limit up-limit down rules. Unpublished working paper, 2020.

K. A. Kim and S. G. Rhee. Price limits performance: Evidence from the Tokyo Stock Exchange. *Journal of Finance*, 52(2):885–901, 1997.

L. E. Kodres and D. P. O'Brien. The existence of Pareto-superior price limits. *American Economic Review*, 84(4):919–932, 1994.

A. S. Kyle. Trading halts and price limits. *Review of Futures Markets*, 7(3):426–434, 1988.

C. M. C. Lee, M. J. Ready, and P. J. Seguin. Volume, volatility, and New York Stock Exchange trading halts. *Journal of Finance*, 49(1):183–214, 1994.

David S. Lee, Justin McCrary, Marcelo J. Moreira, and Jack Porter. Valid t-ratio inference for IV. *American Economic Review*, 112(10):3260–3290, 2022.

Dong Lou, Christopher Polk, and Spyros Skouras. A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134(1):192–213, 2019.

A. Madhavan. Trading mechanisms in securities markets. *Journal of Finance*, 47(2):607–641, 1992.

Ananth Madhavan and Venkatesh Panchapagesan. Price discovery in auction markets: A look inside the black box. *Review of Financial Studies*, 13(3):627–658, 2000.

C. E. Moise and P. Flaherty. "Limit Up-Limit Down" pilot plan and associated events. White paper, U.S. Securities and Exchange Commission, Division of Economic and Risk Analysis, 2017.

W. Ohta. Quote updates and price discovery before market opening. *Toukei Suri (Proceedings of the Institute of Statistical Mathematics)*, 59(1):67–87, 2011. (In Japanese).

Michael S. Pagano and Robert A. Schwartz. A closing call's impact on market quality at Euronext Paris. *Journal of Financial Economics*, 68(3):439–484, 2003.

E. Sanderson and F. Windmeijer. A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*, 190(2):212–221, 2016.

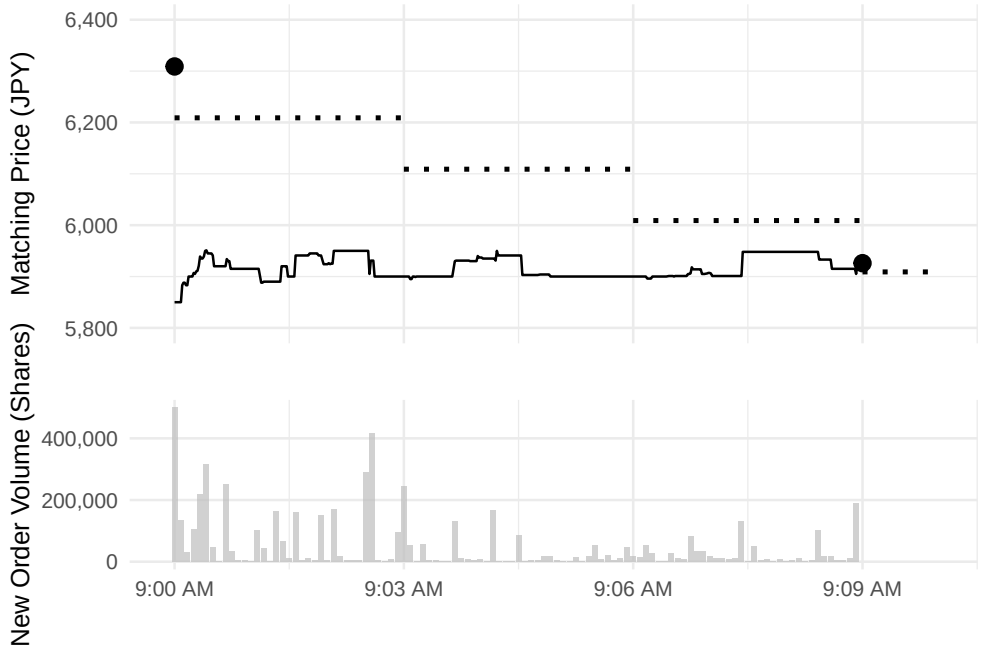
A. Subrahmanyam. Circuit breakers and market volatility: A theoretical perspective. *Journal of Finance*, 49(1):237–254, 1994.

X. Xiao and R. Yamamoto. Price discovery, order submission, and tick size reduction in the preopen period: Evidence from the Tokyo Stock Exchange. *Pacific-Basin Finance Journal*, 63:101428, 2020.

K. Zimmermann. Price discovery in European volatility interruptions. Unpublished Working Paper, 2014.

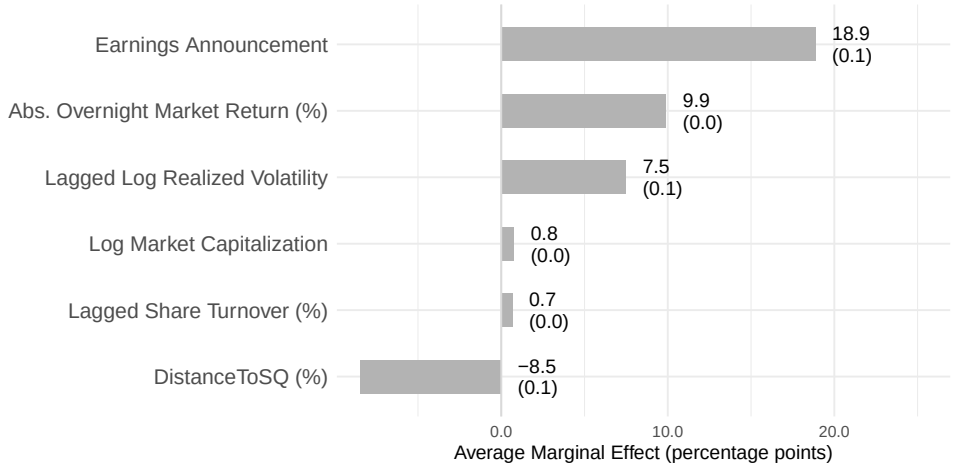
Figures

Figure 1: Trading Pauses at the TSE



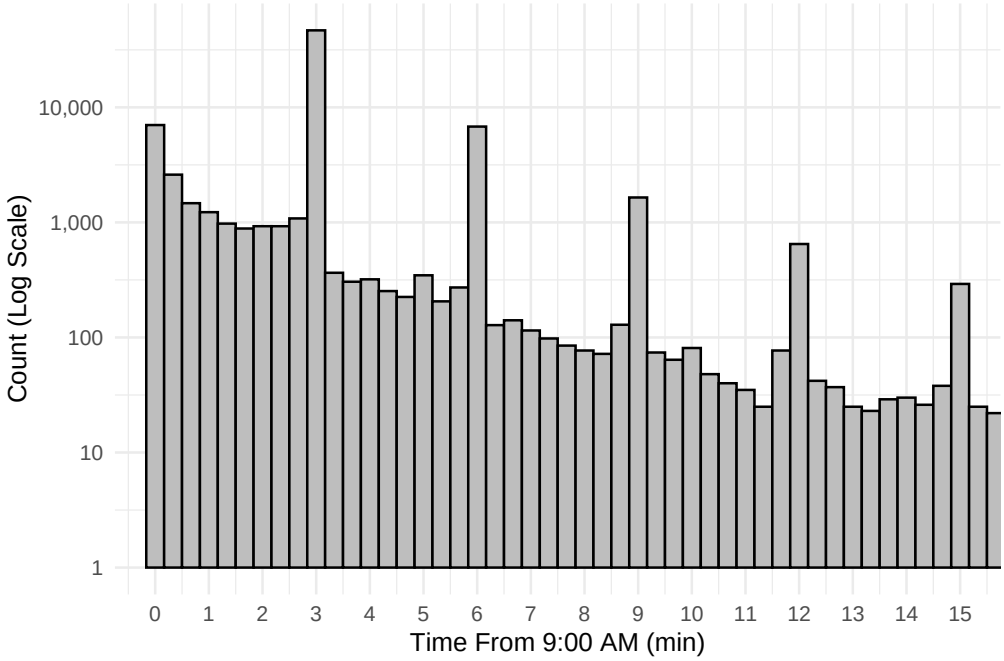
Note: This figure illustrates trading pauses on the Tokyo Stock Exchange (TSE), using as an example Toyota Motor Corporation stock on March 13, 2020, when market volatility surged due to the COVID-19 pandemic. The black points mark the previous day's closing price of 6,309 JPY (at 9:00 AM) and the opening price of 5,926 JPY (at 9:09 AM). The matching price at 9:00 AM was 5,850 JPY—well below the price-limit range—which triggered a trading pause. As a result, the market opening was delayed, and trading began at 9:09 AM at 5,926 JPY, when the price-limit range had been widened for the third time. The solid line indicates the matching price at one-second intervals, while the dotted segments show the special quote, the lower boundary of the price-limit range in this sell-side pause, which was updated every three minutes. The gray bars at the bottom show the volume of new orders at five-second intervals.

Figure 2: Average Marginal Effects From a Probit Model of Trading Pause Probability



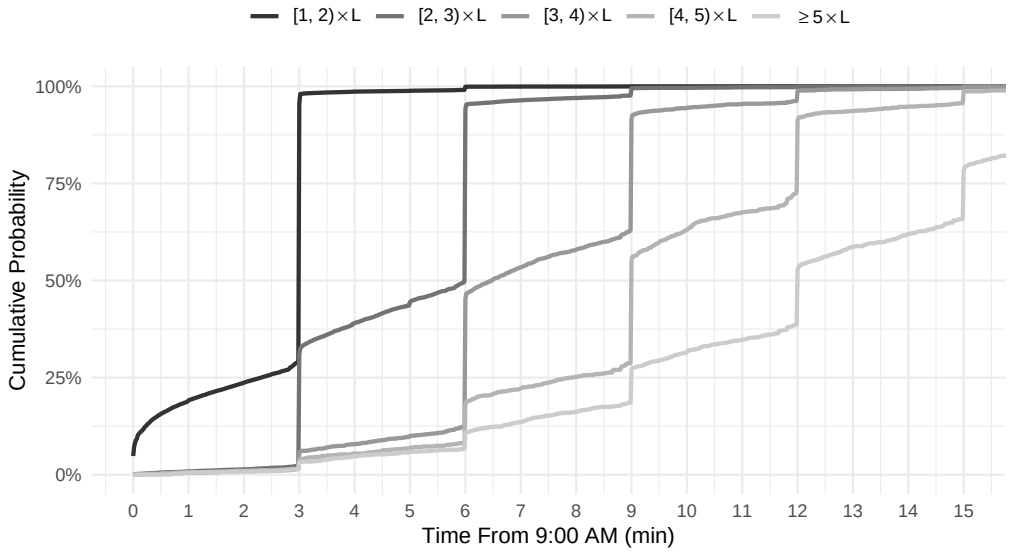
Note: This figure reports average marginal effects from a probit model of trading pause probability at the morning open. Earnings Announcement equals one if the firm released an earnings announcement or forecast revision at or after 3:00 PM on the preceding trading day. Abs. Overnight Market Return is the absolute value of the return from the previous close to the current open on all stocks and J-REITs listed on the TSE, weighted by previous-day market capitalization. Lagged Log Realized Volatility is intraday volatility computed from five-minute sampling over the full trading day on the previous day, expressed in logarithmic form. Lagged Share Turnover is the ratio of trading volume to shares outstanding on the previous day. Log Market Capitalization is the log of market capitalization at the previous close. *DistanceToSQ* (short for Distance to Special Quote) is defined as the percentage distance from the reference price to the triggering threshold of a trading pause; see Section 4.1 for details. Average marginal effects on the probability of a trading pause are expressed in percentage points; standard errors, in the same units, are in parentheses.

Figure 3: Duration of Trading Pauses at the Morning Open



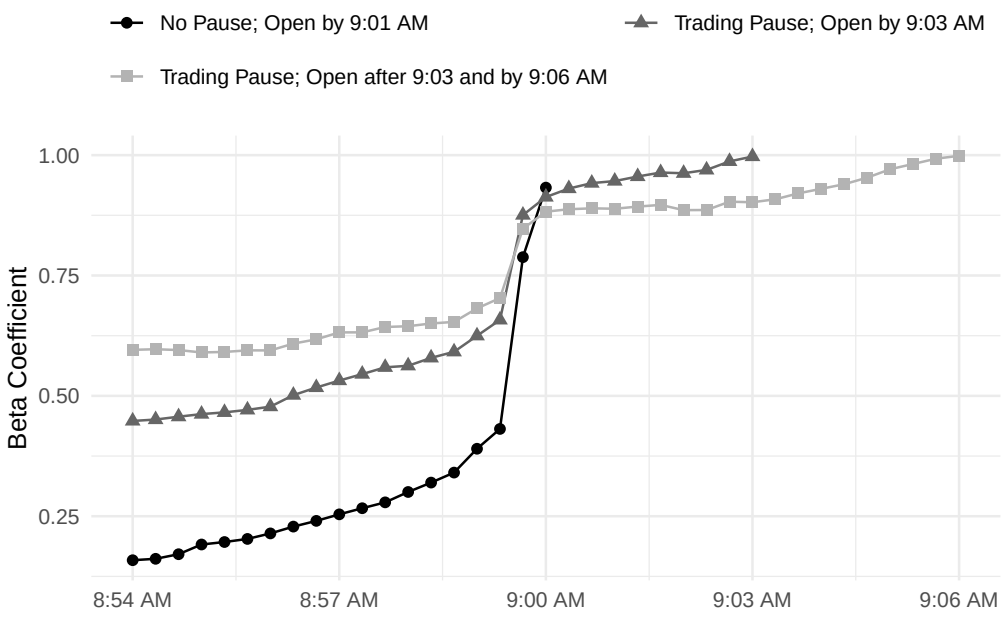
Note: This figure shows the distribution of trading pause durations at the morning open. Counts are plotted on a base-10 logarithmic scale; durations beyond the displayed horizontal range are not shown. The clustering around three-minute intervals (e.g., 3, 6, 9 minutes) corresponds to the scheduled updates of the price-limit range.

Figure 4: Cumulative Distribution of Trading Pause Durations by Order Imbalance



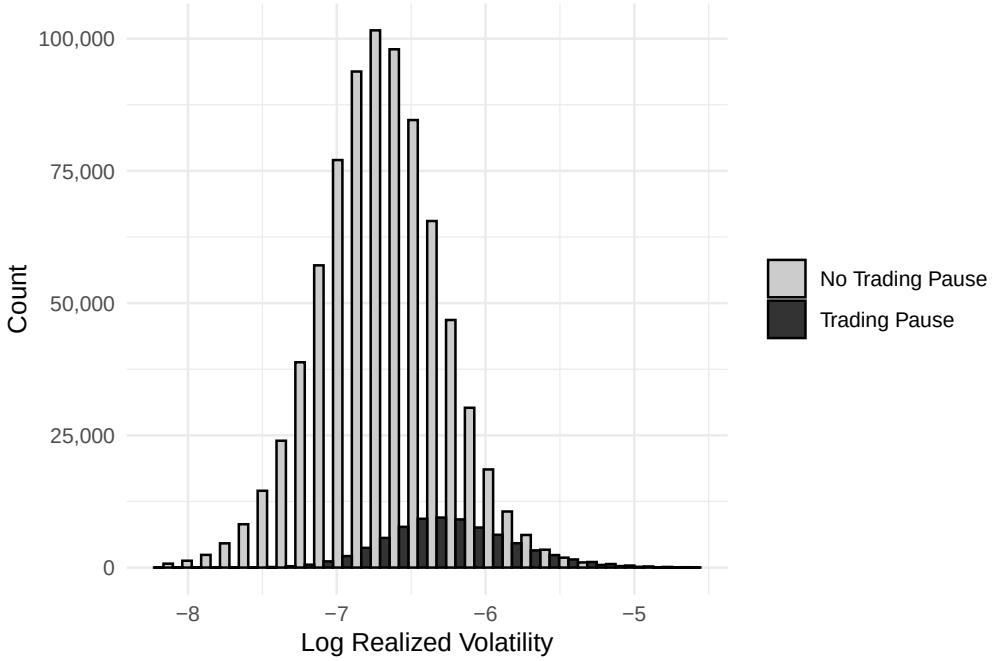
Note: This figure shows the cumulative distribution of the durations of trading pauses triggered at the morning open, stratified by order imbalance. We classify each trading pause by the absolute difference between the matching price at 9:00 AM and the reference price, divided by the range step L. The groups correspond to multiples of L: [1,2), [2,3), [3,4), [4,5), and ≥ 5. Lower multiples (e.g., [1,2) $\times L$) indicate relatively low order imbalance, while higher multiples (e.g., $\geq 5 \times L$) indicate high order imbalance. Pauses without a matching price at 9:00 AM are excluded.

Figure 5: Price Discovery During the Pre-opening Phase With and Without Trading Pauses



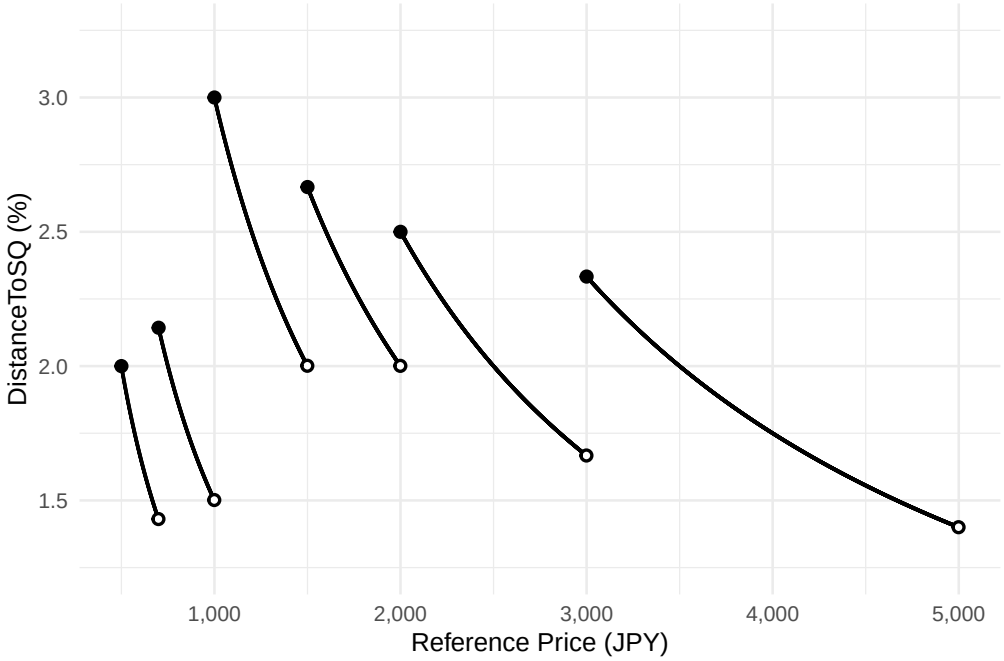
Note: This figure illustrates price discovery during the pre-opening phase across subsamples defined by trading pause occurrence and duration. The "No Pause; Open by 9:01 AM" subsample includes stocks without trading pauses that open by 9:01 AM, while the "Trading Pause; Open by 9:03 AM" subsample includes stocks with trading pauses that end by 9:03 AM, and the "Trading Pause; Open after 9:03 and by 9:06 AM" subsample includes stocks with trading pauses that end after 9:03 AM and by 9:06 AM; the 9:03 and 9:06 AM cutoffs allow a tolerance of one second. For each subsample, we estimate Equation (1) every 20 seconds using matching prices derived from order book data. Beta is the slope of the close-to-close return on the close-to-current-time return, with the latter computed using the matching price.

Figure 6: Realized Volatility With and Without Trading Pauses



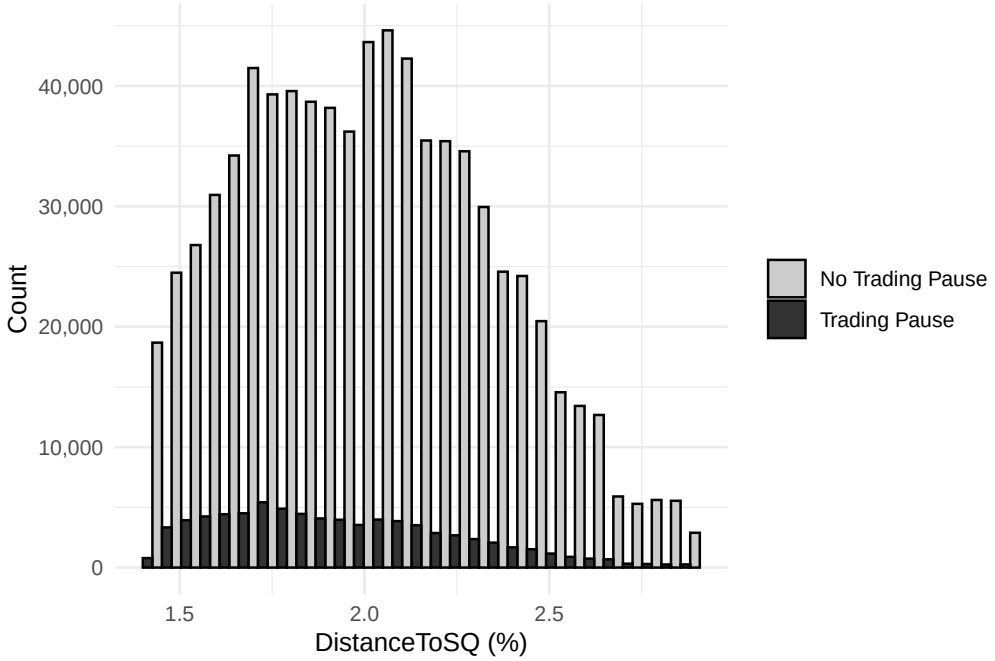
Note: This figure shows the distributions of log realized volatility, calculated using minute-by-minute bid–ask midpoint prices from the start of trading to 30 minutes afterward, conditional on the occurrence of trading pauses. Observations with zero realized volatility are excluded.

Figure 7: Definition of *DistanceToSQ* as a Function of the Reference Price



Note: This figure illustrates *DistanceToSQ* as a function of the reference price. The reference price is the base price at the morning open, normally the previous day's final price or final quote, adjusted for dividend payments and stock splits. *DistanceToSQ* is defined as the percentage distance from the reference price to the triggering threshold of a trading pause.

Figure 8: Distribution of *DistanceToSQ* With and Without Trading Pauses



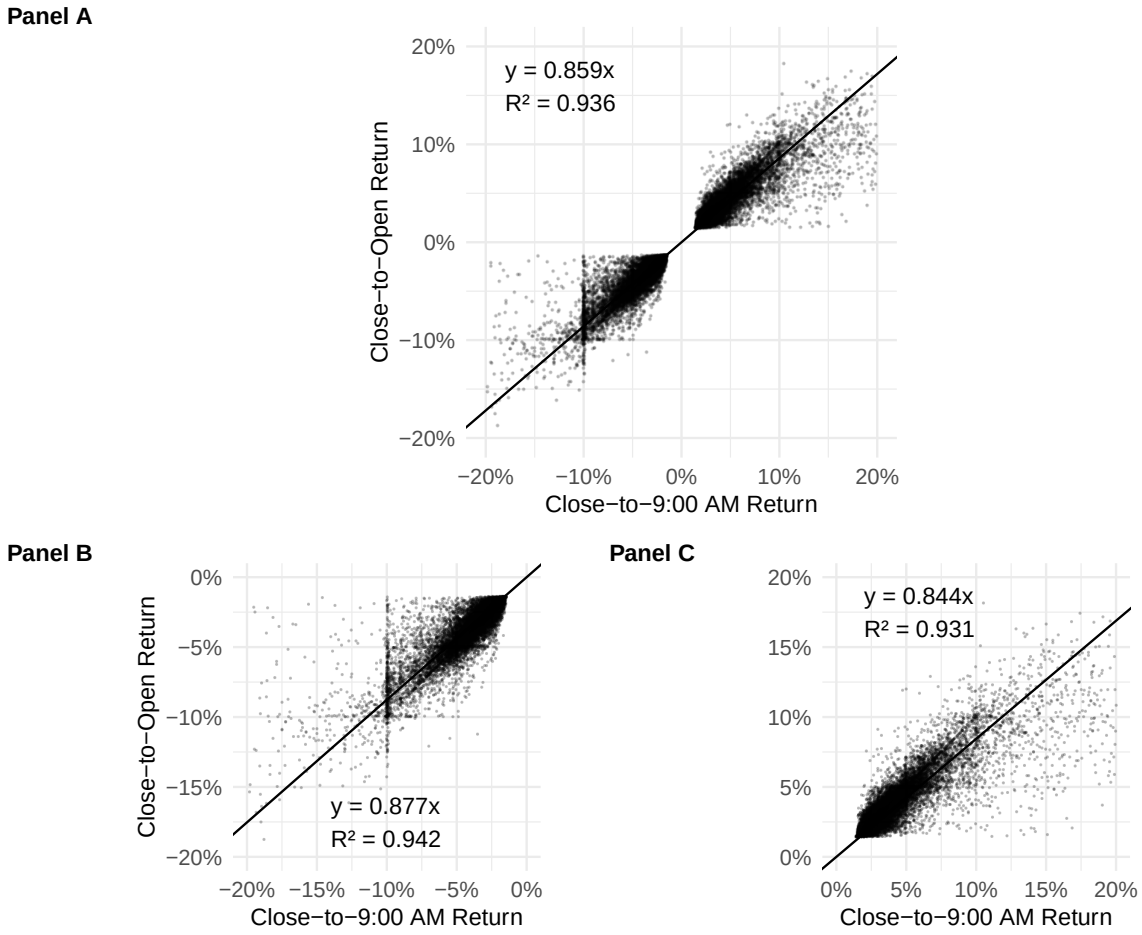
Note: This figure shows the distributions of *DistanceToSQ*, conditional on the presence or absence of trading pauses.

Figure 9: Post-Open Volatility Across 30-Minute Windows



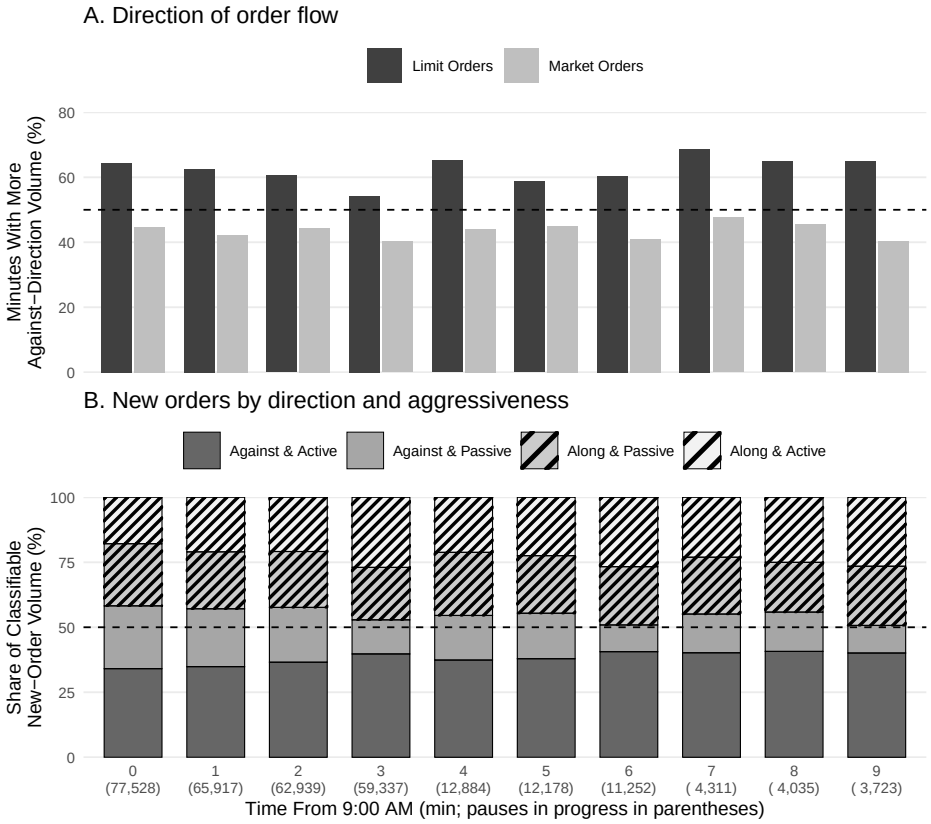
Note: This figure shows IV estimates of the trading pause coefficient on post-open volatility across 30-minute windows that start later in the trading session. Each point represents the coefficient from an IV regression where the dependent variable is the log of minute-by-minute midpoint volatility computed over the $[x, x+30]$-minute window after trading begins. The horizontal axis indicates the window start minute x. The shaded area represents the 95% confidence interval. Control variables follow Table 5, Column (3). All regressions include firm and date fixed effects. Standard errors are two-way clustered by firm and date.

Figure 10: Price Reversals During Trading Pauses



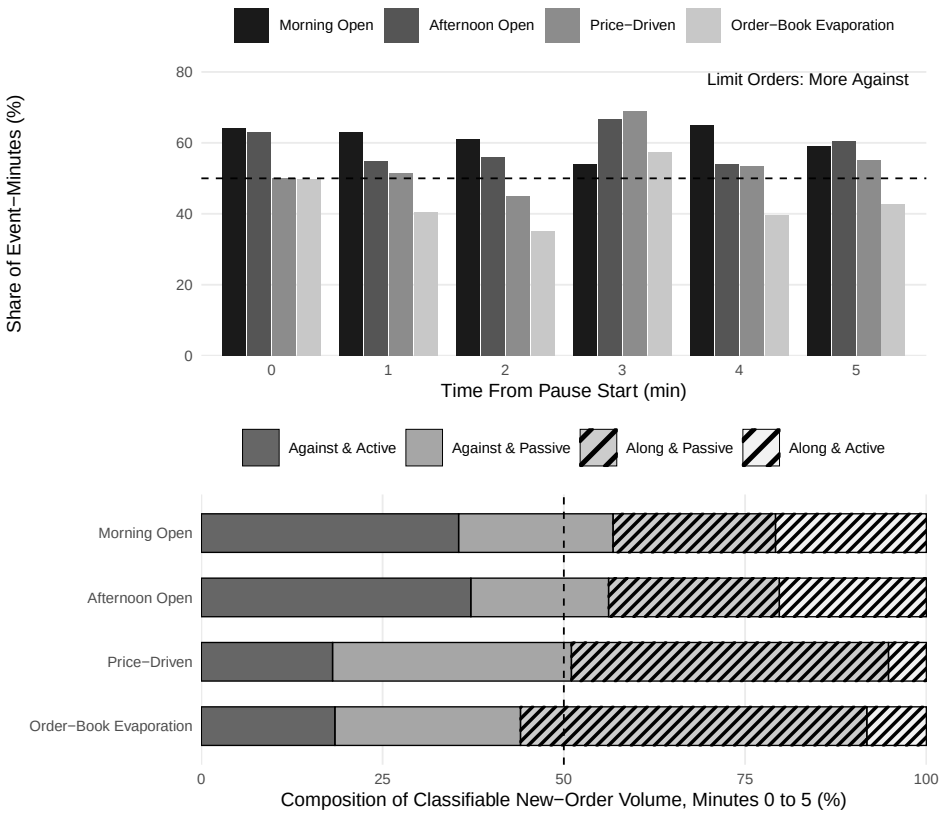
Note: This figure shows price reversals during trading pauses at the morning open. Panel A plots the close-to-open return against the close-to-9:00 AM return (calculated using matching prices) for all observations with trading pauses. The solid line is a regression line with the intercept constrained to zero. Panel B shows the subsample with negative close-to-open returns; Panel C shows the subsample with positive close-to-open returns.

Figure 11: Order Flow During Morning-Open Pauses



Note: This figure summarizes order flow in each minute from 9:00 AM during morning-open pauses of the top-500 stocks; the number of pauses still in progress at the start of each minute is shown in parentheses, and later minutes include only these pauses. The pause direction is buy-side or sell-side according to the special quote. Panel A reports, separately for limit and market orders, the share of pause–minutes in which signed order volume against the pause direction exceeds that along it; order volumes are aggregated within the minute, cancellations enter with negative volume, and pause–minutes with equal volumes in the two directions are excluded. Panel B reports the composition of new-order volume by direction and aggressiveness: an order is active if it is a market order or a limit order priced more aggressively than the matching price (a buy above it or a sell below it), and passive otherwise. Limit orders submitted when no matching price is displayed cannot be classified and are excluded; shares are computed within each pause–minute and averaged with equal weight on every pause. Dashed lines mark 50%.

Figure 12: Order Flow During Trading Pauses by Session and Type



Note: This figure compares order flow during pauses that interrupt continuous trading or extend the afternoon open with order flow during morning-open pauses. The upper panel plots, for each minute from the pause start, the share of event–minutes in which signed limit-order volume against the pause direction exceeds that along it (Panel A of Figure 11); later minutes include only pauses still in progress. The lower panel decomposes classifiable new-order volume in minutes 0–5 by direction and aggressiveness as in Panel B of Figure 11, with equal weight on every event–minute; unclassifiable orders, which lack a matching price, are on average less than 1% of new-order volume at the two openings, 9% in price-driven pauses, and 42% in order-book-evaporation pauses. Dashed lines mark 50%. For continuous-session pauses, the triggering order is excluded. Morning-open series use the sample of Figure 11; the others use the estimation samples of Table 15.

Tables

Table 1: Range-Step Schedule for Trading Pauses

Reference Price P^{ref} (JPY)	Range Step L (JPY)
$P^{ref} < 200$	5
$200 \leq P^{ref} < 500$	8
$500 \leq P^{ref} < 700$	10
$700 \leq P^{ref} < 1{,}000$	15
$1{,}000 \leq P^{ref} < 1{,}500$	30
$1{,}500 \leq P^{ref} < 2{,}000$	40
$2{,}000 \leq P^{ref} < 3{,}000$	50
$3{,}000 \leq P^{ref} < 5{,}000$	70

Note: This table reports the range-step schedule for trading pauses, i.e., the range step L by reference-price level, displaying only reference prices less than 5,000 JPY to save space. During continuous trading the reference price is generally the last execution price; at the morning open it is the base price, normally the previous day's final price or final quote, adjusted for dividend payments and stock splits. These parameters apply to all trading sessions; since November 5, 2019, the closing auction may execute within two range steps of the reference price.

Table 2: Comparison of Single-Stock Trading Pause Rules Across Major Exchanges

Exchange and Mechanism	**Triggering Criteria**	**Duration and Extensions**	**Applicability at Market Open**
TSE (Japan); Trading Pause (Special Quote)	Fixed yen threshold relative to the reference price (the last execution price, or the base price at the morning open)	Can end at any time; the price-limit range widens every three minutes	Applicable
NYSE, Nasdaq (U.S.); LULD Plan	Percentage price bands around the average price over the last five minutes; pause if a limit state is not resolved within 15 seconds	Five minutes, extendable by the listing exchange	Not applicable to the opening auction
LSE (U.K.); Automatic Execution Suspension Period	Percentage thresholds relative to the last auction price and the last execution price	Five minutes with a random period of up to 30 seconds, extendable	Applicable (the opening auction is extended by a price monitoring extension)

Note: This table compares single-stock trading pause rules at the TSE (Japan), NYSE and Nasdaq (U.S.), and LSE (U.K.), focusing on triggering conditions, pause duration, and applicability at market open. Rules are described as of our 2014–2021 sample period; thresholds and extension rules on the LSE vary across securities and trading services.

Table 3: Intraday Frequency of Trading Pauses

	Morning Open	Morning Session	Afternoon Open	Afternoon Session
Total	77,528	5,458	2,208	4,846
	(8.9%)	(0.6%)	(0.3%)	(0.6%)
Sell-Side Pause	38,523	1,700	1,525	1,462
	(4.4%)	(0.2%)	(0.2%)	(0.2%)
Buy-Side Pause	39,005	3,758	683	3,384
	(4.5%)	(0.4%)	(0.1%)	(0.4%)

Note: This table presents the intraday frequency of trading pauses. We count newly occurring trading pauses at the morning open, during the morning continuous session, at the afternoon open, and during the afternoon continuous session. A pause is counted at the morning or afternoon open if its first record is the opening call of the session; a pause that begins after the stock has started trading is counted in the continuous session. Values in parentheses represent the frequency as a proportion of all stock–date observations in the sample. The first row reports the overall frequency, while the subsequent rows distinguish sell-side and buy-side pauses, that is, pauses whose special quote is displayed as an offer (sell-side imbalance) or as a bid (buy-side imbalance).

Table 4: Summary Statistics for Trading Pauses at the Morning Open and Key Variables of Interest

	Mean	Median	Std. Dev.	Min	Max	P5	P95	N
Daily Frequency (%)	8.9	3.8	14.5	0.0	98.2	0.8	36.9	1,763
Duration (min)	3.5	3.0	8.2	0.0	360.0	0.0	8.7	77,528
DistanceToSQ (%)	2.0	2.0	0.4	1.4	6.3	1.5	2.7	872,773
Relative Tick Size (bp)	9.8	7.4	8.4	1.0	125.0	1.7	21.5	872,732
Realized Volatility (bp)	14.4	12.9	7.2	0.0	208.1	6.4	26.9	872,379
Close-to-Open Return (%)	0.0	0.0	1.3	−31.0	29.1	−1.8	1.9	872,732
Close-to-Close Return (%)	0.0	0.0	2.0	−27.6	30.0	−3.0	3.1	872,773
Quoted Spread (bp)	16.1	14.7	9.8	1.4	773.4	4.6	32.6	872,393
Traded Value (Million JPY)	637	231	1,670	0	132,261	33	2,278	872,342
Market Depth Value (Million JPY)	265	162	739	5	94,912	35	704	872,393
Market Capitalization (Billion JPY)	1,036	482	1,737	71	29,877	175	3,889	872,323

Note: This table reports summary statistics for trading pauses at the morning open and key variables of interest. The first two rows present statistics for trading pause characteristics, while the remaining rows cover the full sample. Daily Frequency is the proportion of stocks experiencing trading pauses each day. Duration is the time in minutes from 9:00 AM to the end of the trading pause. *DistanceToSQ* is defined in Section 4.1. Relative Tick Size is the minimum tick size relative to the opening price. Realized Volatility is calculated using minute-by-minute bid–ask midpoint prices from the start of trading to 30 minutes afterward. Quoted Spread is the average quoted bid–ask spread, Traded Value is total traded value, and Market Depth Value is the average value of the limit orders on the ten best price levels on each side of the book (in million JPY, before the tick-size adjustment used in Table A.5)—all computed over the same 30-minute window. Market capitalization is measured using the previous day's closing price.

Table 5: Regression Results for Post-Open Volatility

	OLS			IV		
	(1)	(2)	(3)	(1)	(2)	(3)
Trading Pause	0.199***	0.200***	0.153***	−0.502***	−0.388***	−0.311***
	(0.007)	(0.007)	(0.006)	(0.091)	(0.093)	(0.051)
Relative Tick Size (%)		0.316***	0.190***		0.245***	0.134***
		(0.047)	(0.031)		(0.055)	(0.034)
Log Reference Price			0.010**			0.014***
			(0.005)			(0.005)
Log Market Capitalization			0.013			0.008
			(0.008)			(0.009)
Lagged Share Turnover (%)			0.031***			0.036***
			(0.005)			(0.006)
Lagged Log Realized Volatility			0.454***			0.487***
			(0.007)			(0.008)
Firm & Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.469	0.470	0.554			
First-Stage F-Statistic				622.3	606.2	672.3
Observations	870,177	870,177	869,299	870,177	870,177	869,299

Note: This table reports regression results for Equation (3) with post-open volatility as the dependent variable. The dependent variable is the log of realized volatility, calculated using minute-by-minute bid–ask midpoint prices from the start of trading to 30 minutes afterward (observations with zero volatility are excluded). For comparison and robustness, we report both OLS and IV results under alternative specifications. The IV specifications instrument Trading Pause with *DistanceToSQ*; Table A.2 reports the first-stage regressions. Relative Tick Size is the minimum tick size relative to the opening price. Log Reference Price is the log of the base price at the morning open, normally the previous day's final price or final quote, adjusted for dividend payments and stock splits. Log Market Capitalization is the log of market capitalization at the previous close. Lagged Share Turnover is the ratio of trading volume to shares outstanding on the previous day. Lagged Log Realized Volatility is the log of the previous trading day's realized volatility, $\sqrt{30\,(s_{\mathrm{AM}}^2 + s_{\mathrm{PM}}^2)}$, where s is the standard deviation of five-minute midpoint returns within each session; midpoints use quotes observed outside trading pauses, with the best ask above the best bid, and are carried forward within the session. All specifications include firm and date fixed effects. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Robustness Checks for the Volatility Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Trading Pause	−0.241**	−0.253***	−5.541***	−0.253***	−0.176***	−0.213***	−0.140***	−0.097**
	(0.111)	(0.046)	(0.866)	(0.052)	(0.058)	(0.054)	(0.043)	(0.047)
Relative Tick Size (%)	0.213	0.051	1.438**	1.217***	−0.455***	−0.417***	0.195***	1.372***
	(0.128)	(0.040)	(0.713)	(0.040)	(0.045)	(0.035)	(0.029)	(0.042)
Log Reference Price	0.003	0.016***	0.157**	0.017***	0.024***	0.020***	0.011***	0.002
	(0.010)	(0.006)	(0.074)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)
Log Market Capitalization	−0.010	0.014	−0.079	−0.003	−0.008	0.003	0.009	0.013*
	(0.023)	(0.009)	(0.154)	(0.009)	(0.010)	(0.009)	(0.007)	(0.008)
Lagged Share Turnover (%)	0.102***	0.039***	1.183***	0.024***	0.038***	0.043***	0.033***	0.028***
	(0.016)	(0.006)	(0.173)	(0.004)	(0.007)	(0.007)	(0.005)	(0.004)
Lagged Log Realized Volatility	0.449***	0.484***	6.934***	0.405***	0.528***	0.529***	0.456***	0.387***
	(0.016)	(0.009)	(0.166)	(0.007)	(0.013)	(0.010)	(0.010)	(0.007)
Firm & Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F-Statistic	187.3	678.6	662.4	669.8	662.4	672.4	666.1	664.1
Observations	176,001	811,459	871,091	861,063	871,096	869,387	871,109	871,386

Note: This table presents robustness checks for the IV volatility estimates using alternative samples and volatility measures. Column (1) restricts the sample to TOPIX 100 stocks; Column (2) restricts the sample to observations with reference prices of 500 JPY or higher. Columns (3)–(6) employ alternative volatility measures computed over the first 30 minutes of trading: (3) minute-by-minute midpoint volatility in levels, (4) log transaction-price volatility, (5) log of the largest absolute midpoint return relative to the opening price within the window, and (6) log of the high–low range of midpoint prices, $\ln[\ln(\text{High}/\text{Low})]$. Coefficients in Column (3) are not directly comparable in magnitude due to the scale of the dependent variable. Columns (7) and (8) use same-day volatility over the full trading day, constructed in the same way as the lagged volatility control in Table 5: (7) log midpoint volatility and (8) log transaction-price volatility, for which the price in each five-minute interval is the last trade in that interval. Control variables follow Table 5, Column (3). All specifications include firm and date fixed effects. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Regression Results for Price Efficiency

Panel A: Pause indicator and interaction				
	(1) 9:00 AM		(2) Opening Price	
	OLS	IV	OLS	IV
Trading Pause ($\times 10^3$)	1.122***	−0.758	1.163***	−0.757
	(0.085)	(0.647)	(0.082)	(0.642)
Close-to-9:00 AM Return	0.961***	1.044***		
	(0.002)	(0.022)		
Close-to-Open Return			0.961***	0.985***
			(0.002)	(0.020)
(Interaction Term)	−0.073***	−0.202***	0.020***	−0.020
	(0.004)	(0.033)	(0.004)	(0.032)
Firm & Date FE	No	No	No	No
Adjusted R^2	0.385		0.396	
Excluded-Instrument F		76.8		66.3
Observations	872,029	872,029	872,029	872,029

Panel B: Price efficiency for paused stock–dates			
	(1) 9:00 AM	(2) Opening Price	(3) Encompassing
Close-to-9:00 AM Return	0.888***		−0.010
	(0.003)		(0.016)
Close-to-Open Return		0.981***	0.991***
		(0.003)	(0.017)
p-Value: Close-to-9:00 AM Coefficient = 1	<0.001		
p-Value: Close-to-Open Coefficient = 1		<0.001	0.603
Observations	76,789	76,789	76,789

Note: This table reports regressions for price efficiency. Panel A reports OLS and IV estimates of Equation (4) for the price at 9:00 AM (Column (1); the matching price for paused stocks) and the opening price (Column (2)); the Trading Pause coefficient and its standard error are multiplied by 1,000. Panel B uses the paused stock–dates of the Panel A sample and regresses the close-to-close return on the close-to-9:00 AM return (matching price), the close-to-open return (opening price), and both; its single-return slopes equal $\beta + \gamma$ from the OLS estimates of Panel A, and the comparison is descriptive. Following the unbiasedness regression literature, no fixed effects are included. In the IV columns, the pause indicator and its interaction with the return are instrumented with *DistanceToSQ* and its interaction with the return. Standard errors in parentheses are heteroskedasticity-robust; Excluded-Instrument F is the smaller of the two first-stage Wald statistics for the pause indicator and its interaction, two-way clustered by firm and date; the Sanderson–Windmeijer conditional first-stage F-statistics (Sanderson and Windmeijer, 2016) with the same clustering are at least 130. The p-value rows test whether the return coefficient equals one; in Column (3) of Panel B, the stars on the close-to-9:00 AM coefficient test whether that return adds information. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 8: Regression Results for Post-Open Quoted Spreads

	OLS			IV		
	(1)	(2)	(3)	(1)	(2)	(3)
Trading Pause	−0.654***	−0.249***	−0.356***	−45.948***	−7.136***	−6.249***
	(0.118)	(0.038)	(0.035)	(10.175)	(1.535)	(1.409)
Relative Tick Size (%)		78.683***	78.388***		77.852***	77.687***
		(1.692)	(1.841)		(1.638)	(1.792)
Log Reference Price			0.306**			0.356**
			(0.153)			(0.150)
Log Market Capitalization			−1.138***			−1.198***
			(0.246)			(0.247)
Lagged Share Turnover (%)			−0.788***			−0.720***
			(0.151)			(0.147)
Lagged Log Realized Volatility			1.764***			2.183***
			(0.121)			(0.160)
Firm & Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.678	0.838	0.841			
First-Stage F-Statistic				607.3	593.5	662.4
Observations	872,393	872,393	871,101	872,393	872,393	871,101

Note: This table reports regression results for Equation (3) with post-open quoted spreads as the dependent variable. The dependent variable is the average quoted bid–ask spread in basis points from the start of trading to 30 minutes afterward. For comparison and robustness, we report both OLS and IV results under alternative specifications as in Table 5. All specifications include firm and date fixed effects. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 9: Decomposition of Post-Open Volatility

	Outcome-Specific Samples			Common Sample
Dependent Variable (Log)	Trading Pause	Observations	First-Stage F-Statistic	Trading Pause
Midpoint Realized Volatility, 0–30 min	−0.311***	869,299	672.3	−0.303***
(Benchmark, Table 5)	(0.051)			(0.050)
Path Variation $\sum_m r_m^2$	−0.598***	868,779	675.4	−0.575***
from the Actual Opening Time	(0.102)			(0.101)
Squared Net Return r_{tot}^2	−0.208	836,307	691.7	−0.212
	(0.130)			(0.130)
Variance Ratio $r_{\text{tot}}^2 / \sum_m r_m^2$	0.366***	836,307	691.7	0.363***
	(0.094)			(0.094)

Note: This table reports 2SLS estimates of Equation (3) with the log of each row's quantity as the dependent variable; standard errors are in parentheses below the coefficients. The benchmark row reproduces Table 5, IV Column (3). The other rows use end-of-minute midpoints of one-minute bars starting between each stock–date's actual opening time and 30 minutes later (at least 20 midpoints with the best ask at or above the best bid): r_m is the log return between successive midpoints, $\sum_m r_m^2$ the path variation, $r_{\text{tot}} = \sum_m r_m$ the net return, and $\ln \text{VR} = \ln r_{\text{tot}}^2 - \ln \sum_m r_m^2$ (Section 5.1). The outcome-specific samples differ because the grid-based outcomes have their own data requirements and $\ln r_{\text{tot}}^2$ is undefined for a zero net return. The last column re-estimates all rows on the common sample with all four outcomes observed (835,995 stock–dates; first-stage F-statistic 689.5), on which the variance-ratio coefficient equals the squared-net-return coefficient minus the path-variation coefficient. Magnitudes are not comparable across rows, because the benchmark is a standard deviation. The sample comprises the top-500 stocks, January 2014 to March 2021. All specifications include the controls of Table 5, Column (3), and firm and date fixed effects; standard errors and the first-stage F-statistic are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 10: Order Submission Across Update Cycles During Trading Pauses

Panel A: Arrival timing within update cycles

		After an Update (%)		Before the Next Update (%)	
Order Direction	Submissions	30 s	10 s	30 s	10 s
Against the Pause	1,058,567	18.8	7.6	18.5	7.1
Along the Pause	1,089,766	21.8	10.4	15.1	5.4
Uniform Benchmark	—	16.7	5.6	16.7	5.6

Panel B: Placement of against-direction limit orders

		Position Relative to the Special Quote (%)			
Cycle	Submissions	Short	Near	Beyond	Beyond the Next Update (%)
1	1,567,353	7.2	14.9	77.9	48.2
2	511,256	9.3	18.2	72.5	44.4
3	275,201	9.6	20.4	69.9	42.0
4	127,132	13.2	24.7	62.1	46.0
5	62,910	14.2	29.6	56.1	40.3

Note: This table describes new limit orders submitted during morning-open pauses of the top-500 stocks with at least one update, January 2014 to March 2021 (14,499 pauses). An update cycle runs from the posting or a scheduled update of the special quote to the next scheduled update; only complete cycles within the first 30 minutes of each pause are used, so the final interval, which ends in execution, is excluded. Panel A reports, for cycles after the first (4,810 pauses contribute against-direction orders), the shares of orders arriving within 30 and 10 seconds after an update and before the next update; the Uniform Benchmark row gives the shares under uniform arrival over the 180-second cycle. Panel B classifies against-direction orders in cycles 1–5 by $d = s(P_{\text{order}} - Q_k)/L$, where Q_k is the special quote in cycle k, $s = +1$ (-1) for buy-side (sell-side) pauses, and L is the range step of Table 1: short ($d < -0.5$), near ($|d| \leq 0.5$), beyond ($d > 0.5$), and beyond the next update ($d > 1.5$). Special quotes are reconstructed from the range-step schedule (Section 2) with the initial range step held fixed within a pause; in 8% of pauses the quote crosses a price threshold, where the actual step changes, so later cycles are approximate. Shares are computed over orders. Later cycles are observed only for longer pauses, so differences across cycles are not within-pause changes.

Table 11: Heterogeneous Effects by Size and Liquidity Characteristics

	(1)	(2)	(3)
Trading Pause	−0.341***	−0.255***	−0.263***
	(0.056)	(0.070)	(0.057)
Trading Pause × TOPIX 100	0.106		
	(0.145)		
Trading Pause × Small Stock		−0.127	
		(0.105)	
Trading Pause × Low Turnover			−0.137**
			(0.068)
Controls	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes
First-Stage F-Statistic	639.9	498.7	672.0
Observations	869,299	869,299	869,299

Note: This table reports 2SLS results for Equation (5) examining heterogeneous effects on post-open volatility by size and liquidity characteristics. The dependent variable is defined as in Table 5. We interact the Trading Pause indicator with three indicator variables: TOPIX 100 (equals one for TOPIX 100 constituents), Small Stock (equals one if previous-day market capitalization is below the daily median), and Low Turnover (equals one if previous-day share turnover is below the daily median). For brevity, we omit the coefficients on the standalone indicator variables. Other control variables follow Table 5, Column (3); their coefficients are also omitted. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause; Sanderson–Windmeijer conditional first-stage F-statistics (Sanderson and Windmeijer, 2016), computed with the same clustering, are at least 379 for every endogenous variable in every column. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 12: Heterogeneous Effects by Overnight Return Characteristics

	(1)	(2)	(3)	(4)
Trading Pause	−0.391***	−0.378***	−0.404***	−0.317***
	(0.060)	(0.058)	(0.074)	(0.052)
Trading Pause × Negative Overnight Return	0.165***	0.332***		
	(0.036)	(0.109)		
Trading Pause × Negative Overnight Market Return		−0.197*		
		(0.106)		
Trading Pause × Large Overnight Market Return			0.254***	
			(0.065)	
Trading Pause × Earnings Announcement				0.328**
				(0.151)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
First-Stage F-Statistic	467.4	381.4	484.6	638.5
Observations	869,299	869,299	869,299	869,299

Note: This table reports 2SLS results for Equation (5) examining heterogeneous effects on post-open volatility by overnight return characteristics. The dependent variable is defined as in Table 5. We interact the Trading Pause indicator with four indicator variables: Negative Overnight Return (equals one if the firm's overnight return is negative), Negative Overnight Market Return (equals one if the market overnight return is negative), Large Overnight Market Return (equals one if the absolute market overnight return is above the 90th percentile), and Earnings Announcement (equals one if the firm released an earnings announcement or forecast revision at or after 3:00 PM on the preceding trading day). Negative Overnight Market Return and Large Overnight Market Return are collinear with date fixed effects; only their interactions with Trading Pause are identified. For brevity, we omit the coefficients on the standalone indicator variables. Other control variables follow Table 5, Column (3); their coefficients are also omitted. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause; Sanderson–Windmeijer conditional first-stage F-statistics (Sanderson and Windmeijer, 2016), computed with the same clustering, are at least 234 for every endogenous variable in every column. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 13: Heterogeneous Effects by Lagged Volatility and Return Characteristics

	(1)	(2)	(3)	(4)
Trading Pause	−0.345***	−0.346***	−0.351***	−0.369***
	(0.056)	(0.057)	(0.057)	(0.062)
Trading Pause × High Lagged Volatility	0.150**			
	(0.069)			
Trading Pause × Negative Lagged Return		0.068**	0.026	0.043
		(0.029)	(0.032)	(0.028)
Trading Pause × Large Negative Lagged Return			0.151***	
			(0.052)	
Trading Pause × Negative 3-Day Lagged Return				0.065**
				(0.031)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
First-Stage F-Statistic	579.6	545.6	549.2	491.3
Observations	869,299	869,299	869,299	869,297

Note: This table reports 2SLS results for Equation (5) examining heterogeneous effects on post-open volatility by lagged volatility and return characteristics. The dependent variable is defined as in Table 5. Column (1) interacts Trading Pause with High Lagged Volatility (equals one if previous-day realized volatility is above the 90th percentile). Column (2) interacts Trading Pause with Negative Lagged Return (equals one if previous-day close-to-close return is negative). Column (3) adds an interaction with Large Negative Lagged Return (equals one if the firm's previous-day close-to-close return is below the 10th percentile of previous-day returns pooled across all stock–dates). Column (4) replaces the large-negative-return interaction with Negative 3-Day Lagged Return (equals one if the cumulative close-to-close return over the three previous trading days is negative). For brevity, we omit the coefficients on the standalone indicator variables. Other control variables follow Table 5, Column (3); their coefficients are also omitted. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause; Sanderson–Windmeijer conditional first-stage F-statistics (Sanderson and Windmeijer, 2016), computed with the same clustering, are at least 452 for every endogenous variable in every column. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 14: Spillovers to Non-paused Stocks

	OLS		IV: Peers' Median *DistanceToSQ*		IV: Bottom-Decile Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Peer Ratio (+10 pp)	0.017***	0.010***	−0.016	0.058	−0.099	−0.067
	(0.003)	(0.003)	(0.072)	(0.100)	(0.085)	(0.069)
	[0.001]	[0.001]	[0.038]	[0.069]	[0.041]	[0.051]
AR 95% Confidence Set	—	—	[−0.17, +0.12]	[−0.15, +0.27]	[−0.31, +0.05]	[−0.22, +0.07]
First Stage (pp)	—	—	−0.426***	−0.227***	0.190***	0.136***
	—	—	(0.044)	(0.034)	(0.033)	(0.022)
First-Stage F-Statistic	—	—	91.9	44.2	33.8	38.8
Reduced Form (%)	—	—	0.067	−0.131	−0.188	−0.091
	—	—	(0.306)	(0.229)	(0.146)	(0.093)
Own-Stock Controls	Yes	Yes	Yes	Yes	Yes	Yes
Peer Controls	No	Yes	No	Yes	No	Yes
Firm & Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations				3,062,230		

Note: This table reports OLS and IV estimates of Equation (3) for non-paused stock–dates of First Section stocks; J-REITs are excluded. The dependent variable is the log 30-minute realized volatility of Table 5. Peer Ratio is the leave-one-out share of same-industry peers, based on the TSE's 33-industry classification, paused at the morning open; coefficients, standard errors, and confidence sets are per 10-percentage-point increase. The instruments are the industry median of peers' *DistanceToSQ* (Columns (3)–(4)) and the share of peers whose *DistanceToSQ* is below the pooled tenth percentile (Columns (5)–(6)). First-stage (percentage points of Peer Ratio) and reduced-form (percent) coefficients are per one standard deviation of the instrument. All specifications include the own-stock controls of Table 5, Column (3), and firm and date fixed effects; peer controls are leave-one-out industry–date means of peers' natural-spline basis of the log base price (eight degrees of freedom), relative tick size, log market capitalization, turnover, log lagged volatility, one- and three-day lagged returns, and absolute lagged return, together with log industry size. Standard errors in parentheses are two-way clustered by industry and date, and those in brackets by firm and date; first-stage F-statistics, Anderson–Rubin 95% confidence sets, and stars use the industry–date clustering. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 15: Post-Event Volatility by Session and Pause Type

	Morning Open	Afternoon	Continuous Session	
	(Section 4)	Open	Price-Driven	Order-Book Evaporation
Trading Pause	-0.311***	-0.500*	0.018	-0.577
	(0.051)	(0.262)	(0.239)	(0.366)
First-Stage F-Statistic	672.3	32.5	30.1	19.6
AR 95% Confidence Set	$[-0.41, -0.21]$	$[-1.20, -0.07]$	$[-0.52, +0.48]$	$[-1.72, -0.01]$
Treated Events	77,150	3,686	1,406	2,032

Note: This table reports 2SLS estimates of the effect of a trading pause on log 30-minute post-event realized volatility by pause type. The morning-open column reproduces Table 5, IV Column (3), with outcomes measured from each stock's opening trade. The other columns use the intraday design of Section 6: each event on a stock–date with a single pause is matched to same-stock control windows at the same clock time on the ten preceding trading days, and the controls of Table 5, Column (3), are augmented with log pre-event realized volatility (from up to 30 preceding one-minute returns within the same session; for afternoon-open events, the last 30 minutes of the morning session). In these columns, outcome windows begin at the treated pause's ending clock time on both event and matched control days, and observations are event-weighted; all specifications include firm and date fixed effects. The instrument is the percentage distance from the ex-ante reference price to the triggering threshold (Appendix B). Continuous-session events last at least one second and afternoon-open events at least 30 seconds; windows that would extend past the 11:30 AM lunch break are excluded. Treated Events counts the paused observations in each column; the other observations are non-paused stock–dates in the morning-open column and matched control windows in the other columns. Standard errors in parentheses and first-stage F-statistics are two-way clustered by firm and date. Stars are based on conventional t-tests; the Anderson–Rubin 95% confidence sets are robust to weak instruments. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

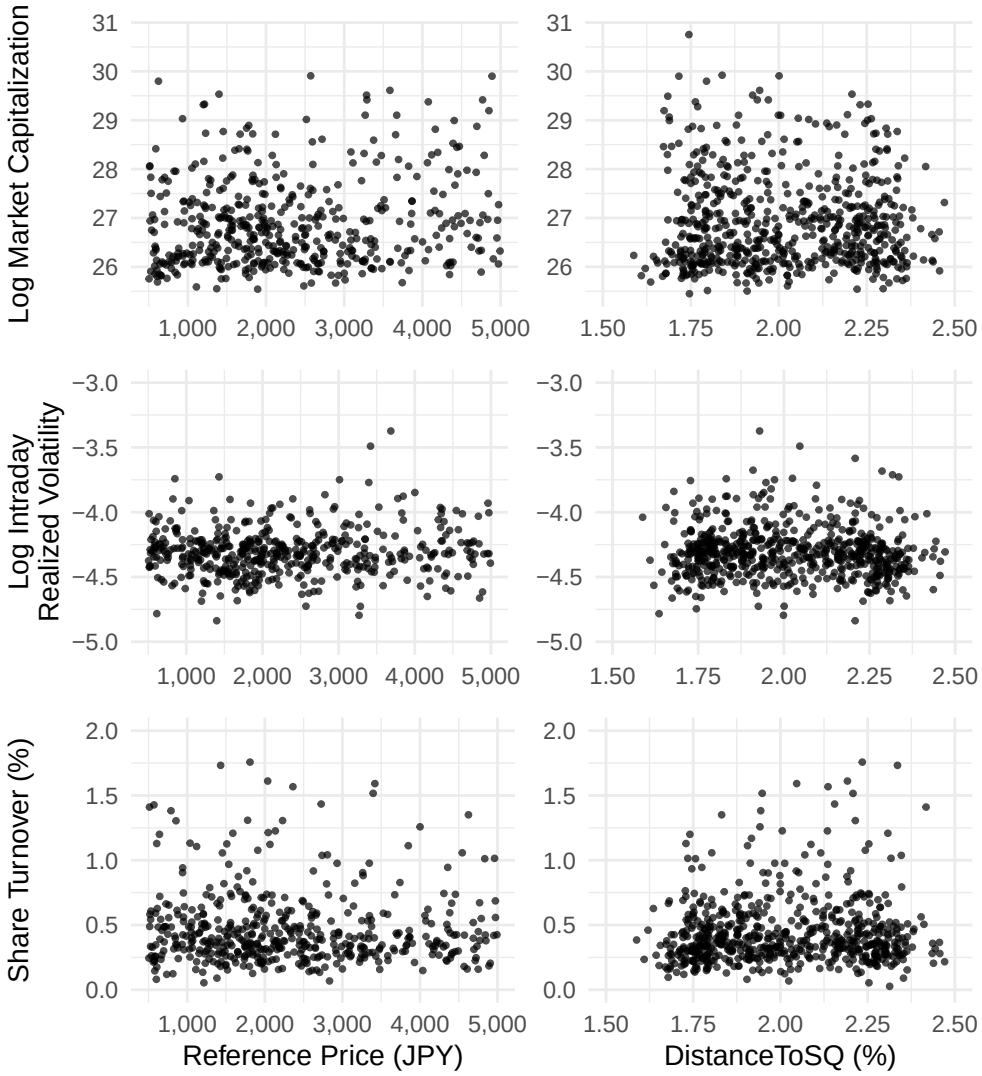
Table 16: Tests for Magnet Effects During Continuous Trading

Panel A: Full First Section universe		
	All Minutes	No Pause in Next 2 min
Relevance (next five minutes)		
Pause Within Five Minutes (bp)	−16.38*** (1.60)	—
Base Rate of a Pause Within Five Minutes (bp)	28.8	—
Trading outcomes (next two minutes)		
Log(1 + Traded Value)	−0.028 (0.038)	−0.027 (0.038)
|Return| (%)	0.005*** (0.001)	0.006*** (0.001)
Continuation in the Direction of the Move	0.013*** (0.001)	0.013*** (0.001)
Aggressor Imbalance Toward the Move	−1.425 (1.084)	−1.346 (1.080)
Observations	5,608,714	5,601,006
Panel B: Results by size group		
	Top-500 Stocks	Outside Top 500
Base Rate of a Pause Within Five Minutes (bp)	13.5	30.9
Pause Within Five Minutes (bp)	−0.49 (3.18)	−18.41*** (1.73)
Log(1 + Traded Value)	−0.109*** (0.041)	−0.021 (0.042)
Continuation in the Direction of the Move	0.018*** (0.004)	0.012*** (0.001)
Observations	670,684	4,938,028

Note: This table reports regressions of trading behavior in the continuous session on $z = L(p)/p \times 100$, *DistanceToSQ* of Equation (2) evaluated at the current price p, where $L(p)$ is the range step of Table 1. Each cell is the coefficient on z from a separate regression, with its standard error in parentheses; pause-probability entries and base rates are in basis points of probability. The relevance row uses an indicator for a pause starting within the next five minutes; the other outcomes are measured over the next two minutes. Continuation indicates that the next two-minute return has the sign of the preceding move; aggressor imbalance is the signed trade imbalance toward the move, scaled by preceding activity and winsorized at 1%/99%. Observations are stock–minutes of the full First Section universe in the continuous session (9:35 AM–11:20 AM and 12:40 PM–2:45 PM, 2014–2021) with $p \geq 100$ JPY, an absolute five-minute return of at least 0.8% (computed from last transaction prices, with the earlier price at most 15 minutes old), and a trade in the current minute, excluding minutes during a pause or within ten minutes after one (capped at 120 minutes after the pause start). Controls are $\log p$, relative tick size, the absolute five-minute return, and $\log(1 + \text{traded value})$ over the preceding five minutes, with firm and date fixed effects; standard errors are two-way clustered by firm and date. The relevance of z is not established in the top-500 group, so its coefficients are not interpreted as responses to anticipation. Panel B regressions drop their own singleton groups. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

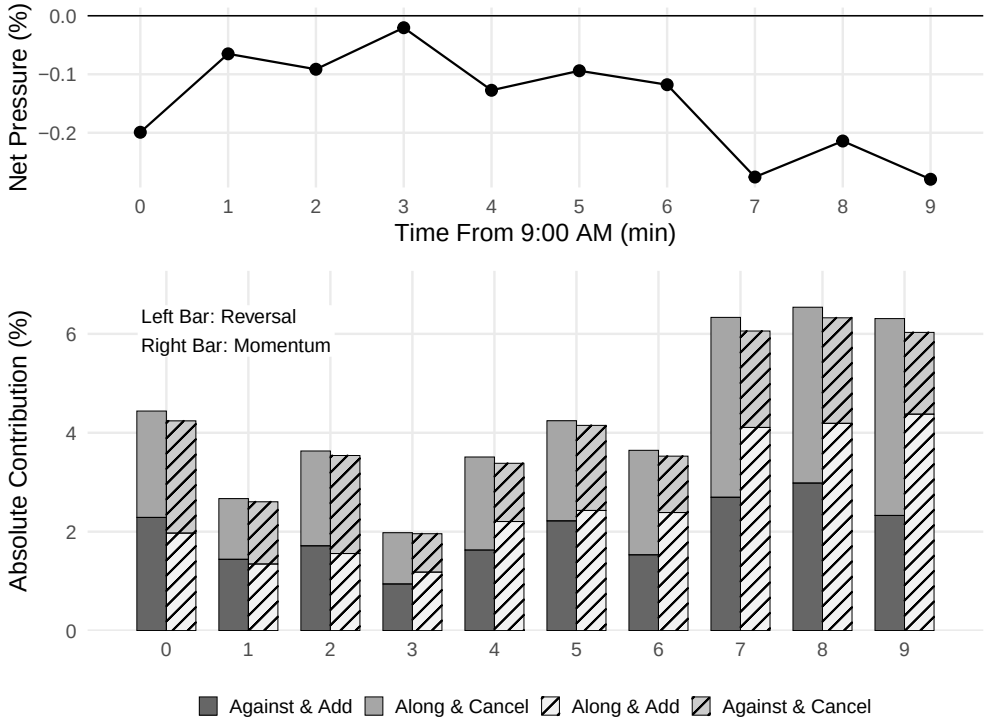
Appendix A. Supplementary Figures and Tables

Figure A.1: Distributions of Market Capitalization, Intraday Realized Volatility, Share Turnover, Reference Price, and *DistanceToSQ*



Note: This figure illustrates the distributions and relationships among market capitalization, intraday realized volatility, share turnover, reference price, and *DistanceToSQ*. Data are averaged by stock code. Market capitalization and intraday realized volatility are shown on a logarithmic scale. Intraday realized volatility is computed using five-minute bid–ask midpoint prices. Share turnover is the ratio of daily trading volume to shares outstanding.

Figure A.2: Net Order Pressure and Its Components During Trading Pauses



Note: This figure plots signed net order pressure during trading pauses, scaled by the stock's mean previous-day trading volume on its pause dates, weighted by the number of order records on each date. The upper panel reports the event-time average across trading-pause events. We first assign an "upward-pressure" sign: bid-side new orders and ask-side cancellations are treated as upward pressure, whereas ask-side new orders and bid-side cancellations are treated as downward pressure. We then sign pressure so that positive values indicate pressure along the pause direction and negative values indicate pressure against it (along and against are defined as in Figure 11). The lower panel decomposes the net measure into four components by order direction (Along vs. Against) and event type (Add vs. Cancel). For comparability, the panel plots absolute contributions: the left bar ("Reversal") stacks Against & Add and Along & Cancel, and the right bar ("Momentum") stacks Along & Add and Against & Cancel. Unlike Figure 11, which weights each pause equally, this figure weights pauses by their number of order records and therefore gives more weight to pauses with heavy order flow; in these pauses, new along-direction orders exceed new against-direction orders from about the fourth minute.

Table A.1: Probit Estimates of Trading Pause Probability

	(1)	(2)	(3)
DistanceToSQ (%)	−0.493***	−0.724***	−0.785***
	(0.006)	(0.007)	(0.007)
Earnings Announcement		1.774***	1.753***
		(0.013)	(0.013)
Abs. Overnight Market Return (%)		0.997***	0.920***
		(0.003)	(0.003)
Lagged Log Realized Volatility			0.696***
			(0.006)
Lagged Share Turnover (%)			0.063***
			(0.003)
Log Market Capitalization			0.072***
			(0.002)
Firm & Date FE	No	No	No
Log Likelihood	−257,318.9	−183,906.9	−174,926.9
Pseudo R^2 (McFadden)	0.017	0.297	0.331
Observations	872,773	872,773	871,468

Note: This table reports probit coefficient estimates for the probability of a trading pause, the model underlying Figure 2. See Figure 2 for variable definitions; Figure 2 reports average marginal effects from Column (3). Standard errors are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.2: First-Stage Regression: Trading Pause Indicator on *DistanceToSQ*

	(1)	(2)	(3)
DistanceToSQ (%)	−0.082***	−0.082***	−0.084***
	(0.003)	(0.003)	(0.003)
Relative Tick Size (%)		−0.006	−0.027**
		(0.014)	(0.013)
Log Reference Price			0.001
			(0.002)
Log Market Capitalization			−0.014***
			(0.004)
Lagged Share Turnover (%)			0.011***
			(0.002)
Lagged Log Realized Volatility			0.074***
			(0.003)
Firm & Date FE	Yes	Yes	Yes
Adjusted R^2	0.285	0.285	0.291
Partial R^2	0.010	0.010	0.010
Excluded-Instrument F	622.3	606.2	672.3
Observations	870,177	870,177	869,299

Note: This table reports first-stage regressions for the 2SLS estimation in Table 5. The dependent variable is the trading pause indicator. The instrument is *DistanceToSQ*. All specifications include firm and date fixed effects. Columns (1)–(3) use the control sets and samples of the corresponding columns of Table 5. Partial R^2 measures the incremental explanatory power of *DistanceToSQ* relative to the restricted model without the instrument. Excluded-Instrument F is the Wald F-statistic for the null that the coefficient on *DistanceToSQ* equals zero. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.3: Reduced-Form Regression: Post-Open Volatility on *DistanceToSQ*

	(1)	(2)	(3)
DistanceToSQ (%)	0.041***	0.032***	0.026***
	(0.007)	(0.007)	(0.004)
Relative Tick Size (%)		0.248***	0.142***
		(0.051)	(0.032)
Log Reference Price			0.014***
			(0.005)
Log Market Capitalization			0.012
			(0.008)
Lagged Share Turnover (%)			0.033***
			(0.005)
Lagged Log Realized Volatility			0.464***
			(0.007)
Firm & Date FE	Yes	Yes	Yes
Adjusted R^2	0.458	0.458	0.547
Observations	870,177	870,177	869,299

Note: This table reports reduced-form regressions for the 2SLS estimation in Table 5. The dependent variable is the log of realized volatility, as defined in Table 5. The main explanatory variable is *DistanceToSQ*. All specifications include firm and date fixed effects. Columns (1)–(3) use the control sets and samples of the corresponding columns of Table 5. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.4: Regression Results for Price Efficiency With Date Fixed Effects

Panel A: Pause indicator and interaction, date fixed effects				
	(1) 9:00 AM		(2) Opening Price	
	OLS	IV	OLS	IV
Trading Pause ($\times 10^3$)	0.523***	−0.219	0.406**	−0.251
	(0.200)	(0.928)	(0.182)	(0.929)
Close-to-9:00 AM Return	0.931***	0.992***		
	(0.007)	(0.032)		
Close-to-Open Return			0.949***	0.952***
			(0.007)	(0.027)
(Interaction Term)	−0.089***	−0.175***	0.001	−0.003
	(0.009)	(0.043)	(0.009)	(0.038)
Fixed Effects	Date	Date	Date	Date
Adjusted R^2	0.550		0.559	
Excluded-Instrument F		110.6		90.3
Observations	872,029	872,029	872,029	872,029

Panel B: Price efficiency for paused stock–dates, date fixed effects			
	(1) 9:00 AM	(2) Opening Price	(3) Encompassing
Close-to-9:00 AM Return	0.843***		0.010
	(0.010)		(0.024)
Close-to-Open Return		0.950***	0.939***
		(0.010)	(0.026)
p-Value: Close-to-9:00 AM Coefficient = 1	<0.001		
p-Value: Close-to-Open Coefficient = 1		<0.001	0.020
Observations	76,789	76,789	76,789

Note: This table re-estimates both panels of Table 7 with date fixed effects and standard errors two-way clustered by firm and date, on the Table 7 sample (872,029 stock–dates in Panel A and the 76,789 paused stock–dates in Panel B). The excluded-instrument F is the smaller of the two first-stage statistics (for the pause indicator and its interaction), as in Table 7. Because both the fixed effects and the standard errors differ from Table 7, differences in p-values reflect both changes. The Trading Pause coefficient and its standard error in Panel A are multiplied by 1,000 for presentation. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively, for the hypothesis of a zero coefficient. The Panel B slopes are $\beta + \gamma$ from the date-fixed-effects interaction model of Panel A, so the date effects are common to paused and non-paused stock–dates. In Panel B, the rows "p-Value: Close-to-9:00 AM Coefficient = 1" and "p-Value: Close-to-Open Coefficient = 1" test whether the corresponding return coefficient equals one; in Column (3), the stars on the close-to-9:00 AM coefficient test whether that return adds information in the encompassing regression.

Table A.5: Regression Results for Post-Open Traded Value, Market Depth, and Opening-Call Traded Value

Panel A: Traded Value and Market Depth						
	Traded Value			Market Depth (Tick-Adjusted)		
	(1)	(2)	(3)	(4)	(5)	(6)
Trading Pause	1.018***	0.776**	0.166	0.683***	0.654***	−0.031
	(0.302)	(0.301)	(0.139)	(0.242)	(0.251)	(0.132)
Relative Tick Size (%)		−0.486**	−0.591***		−0.057	0.094
		(0.215)	(0.115)		(0.162)	(0.122)
Log Reference Price			−0.054***			−0.002
			(0.013)			(0.014)
Log Market Capitalization			0.970***			0.723***
			(0.028)			(0.027)
Lagged Share Turnover (%)			0.345***			0.154***
			(0.057)			(0.024)
Lagged Log Realized Volatility			0.385***			−0.322***
			(0.031)			(0.026)
Firm & Date FE	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F-Statistic	607.4	593.6	662.5	607.3	593.5	662.4
Observations	872,342	872,342	871,056	872,393	872,393	871,101
Panel B: Opening-Call Traded Value						
	(7)	(8)	(9)			
Trading Pause	0.307	0.489*	−0.140			
	(0.278)	(0.273)	(0.120)			
Relative Tick Size (%)		0.368**	0.327***			
		(0.184)	(0.089)			
Log Reference Price			−0.035***			
			(0.010)			
Log Market Capitalization			0.935***			
			(0.024)			
Lagged Share Turnover (%)			0.323***			
			(0.056)			
Lagged Log Realized Volatility			0.304***			
			(0.030)			
Firm & Date FE	Yes	Yes	Yes			
First-Stage F-Statistic	609.5	595.7	664.6			
Observations	872,773	872,732	871,431			

Note: This table reports 2SLS results for Equation (3) under the specifications of Table 5 for three measures of trading activity: the logs of traded value and of tick-adjusted depth over the first 30 minutes of trading, where tick-adjusted depth is Depth Value divided by the relative tick size and Depth Value is the Market Depth Value of Table 4 (the ten best price levels on each side of the book), and the log of the value executed in the opening call auction. Columns (7)–(9) use the first opening-call record of each stock–date (within one minute of the opening time for 99.99% of stock–dates) and require only a positive opening-call execution and the controls of each specification, so their samples are larger; positive lagged realized volatility is required only in Column (9). Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.6: Heterogeneous Effects by Size and Liquidity Characteristics: OLS Estimates

	(1)	(2)	(3)
Trading Pause	0.161***	0.139***	0.152***
	(0.006)	(0.007)	(0.006)
Trading Pause × TOPIX 100	−0.038***		
	(0.008)		
Trading Pause × Small Stock		0.029***	
		(0.006)	
Trading Pause × Low Turnover			0.002
			(0.005)
Controls	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes
Adjusted R^2	0.554	0.554	0.555
Observations	869,299	869,299	869,299

Note: This table reports OLS counterparts to the IV estimates in Table 11. See Table 11 for variable definitions. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.7: Heterogeneous Effects on Quoted Spreads by Size and Liquidity Characteristics

	(1)	(2)	(3)
Trading Pause	−8.998***	−2.251*	−6.108***
	(1.825)	(1.294)	(1.727)
Trading Pause × TOPIX 100	13.654***		
	(2.932)		
Trading Pause × Small Stock		−8.543***	
		(3.048)	
Trading Pause × Low Turnover			−0.097
			(2.606)
Controls	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes
First-Stage F-Statistic	627.8	493.6	663.3
Observations	871,101	871,101	871,101

Note: This table replaces the dependent variable in Table 11 with the average quoted bid–ask spread in basis points and reports 2SLS results. See Table 11 for variable definitions. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.8: Heterogeneous Effects by Overnight Return Characteristics: OLS Estimates

	(1)	(2)	(3)	(4)
Trading Pause	0.166***	0.168***	0.228***	0.103***
	(0.008)	(0.008)	(0.007)	(0.005)
Trading Pause × Negative Overnight Return	−0.029***	−0.019		
	(0.010)	(0.012)		
Trading Pause × Negative Overnight Market Return		−0.014		
		(0.013)		
Trading Pause × Large Overnight Market Return			−0.230***	
			(0.008)	
Trading Pause × Earnings Announcement				0.126***
				(0.013)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.554	0.554	0.558	0.567
Observations	869,299	869,299	869,299	869,299

Note: This table reports OLS counterparts to the IV estimates in Table 12. See Table 12 for variable definitions. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.9: Heterogeneous Effects on Quoted Spreads by Overnight Return Characteristics

	(1)	(2)	(3)	(4)
Trading Pause	−7.649***	−7.263***	−8.885***	−6.293***
	(1.676)	(1.594)	(1.987)	(1.428)
Trading Pause × Negative Overnight Return	2.891***	7.699***		
	(0.766)	(2.235)		
Trading Pause × Negative Overnight Market Return		−5.665***		
		(1.986)		
Trading Pause × Large Overnight Market Return			7.201***	
			(1.651)	
Trading Pause × Earnings Announcement				2.397
				(1.484)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
First-Stage F-Statistic	462.1	378.4	479.1	628.5
Observations	871,101	871,101	871,101	871,101

Note: This table replaces the dependent variable in Table 12 with the average quoted bid–ask spread in basis points and reports 2SLS results. See Table 12 for variable definitions. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.10: Heterogeneous Effects by Lagged Volatility and Return Characteristics: OLS Estimates

	(1)	(2)	(3)	(4)
Trading Pause	0.170***	0.165***	0.168***	0.176***
	(0.007)	(0.007)	(0.007)	(0.008)
Trading Pause × High Lagged Volatility	−0.077***			
	(0.007)			
Trading Pause × Negative Lagged Return		−0.024***	−0.007	−0.012*
		(0.006)	(0.006)	(0.007)
Trading Pause × Large Negative Lagged Return			−0.070***	
			(0.007)	
Trading Pause × Negative 3-Day Lagged Return				−0.034***
				(0.007)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.555	0.554	0.555	0.554
Observations	869,299	869,299	869,299	869,297

Note: This table reports OLS counterparts to the IV estimates in Table 13. See Table 13 for variable definitions. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.11: Heterogeneous Effects on Quoted Spreads by Lagged Volatility and Return Characteristics

	(1)	(2)	(3)	(4)
Trading Pause	−6.597***	−7.271***	−7.279***	−7.759***
	(1.466)	(1.572)	(1.572)	(1.658)
Trading Pause × High Lagged Volatility	1.791			
	(2.174)			
Trading Pause × Negative Lagged Return		1.963***	1.105**	1.437***
		(0.509)	(0.532)	(0.468)
Trading Pause × Large Negative Lagged Return			2.476**	
			(1.076)	
Trading Pause × Negative 3-Day Lagged Return				1.402***
				(0.541)
Controls	Yes	Yes	Yes	Yes
Firm & Date FE	Yes	Yes	Yes	Yes
First-Stage F-Statistic	570.2	539.7	543.1	487.4
Observations	871,101	871,101	871,101	871,099

Note: This table replaces the dependent variable in Table 13 with the average quoted bid–ask spread in basis points and reports 2SLS results. See Table 13 for variable definitions. The first-stage F-statistic tests the coefficient on *DistanceToSQ* in the first stage for Trading Pause. Standard errors in parentheses are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.12: Robustness of the Continuous-Session and Afternoon-Open Estimates

Specification	Trading Pause	First-Stage F-Statistic	AR 95% Confidence Set	Treated Events
Afternoon Open (Duration $\geq$ 30 s)				
Baseline: $k = 10$ Control Days	-0.500* (0.262)	32.5	$[-1.20, -0.07]$	3,686
$k = 5$ Control Days	-0.717** (0.351)	23.7	$[-1.74, -0.17]$	3,520
$k = 3$ Control Days	-0.560 (0.454)	13.1	$[-2.25, +0.11]$	3,267
All Durations ($k = 5$ Control Days)	-0.723*** (0.209)	59.4	$[-1.22, -0.36]$	7,377
Price-Driven (Continuous Session, Duration $\geq$ 1 s)				
Baseline: $k = 10$ Control Days	0.018 (0.239)	30.1	$[-0.52, +0.48]$	1,406
$k = 5$ Control Days	0.156 (0.278)	22.8	$[-0.49, +0.71]$	1,215
$k = 3$ Control Days	-0.106 (0.261)	26.6	$[-0.73, +0.38]$	1,002
Reference-Price Tolerance 0.1 s	-0.005 (0.241)	30.0	$[-0.55, +0.46]$	1,406
Excluding Starts at 12:30–12:59 PM	0.036 (0.245)	27.7	$[-0.52, +0.51]$	1,264
Order-Book Evaporation (Continuous Session, Duration $\geq$ 1 s)				
Baseline: $k = 10$ Control Days	-0.577 (0.366)	19.6	$[-1.72, -0.01]$	2,032
$k = 5$ Control Days	-0.968 (0.698)	7.6	$[-5.02, -0.03]$	1,887
$k = 3$ Control Days	-0.870 (0.870)	4.3	$[-25.72, +0.22]$	1,691
Reference-Price Tolerance 0.1 s	-0.575 (0.368)	19.3	$[-1.73, -0.01]$	2,031
Excluding Starts at 12:30–12:59 PM	-0.445 (0.335)	20.4	$[-1.45, +0.09]$	1,872

Note: This table re-estimates Table 15 under alternative design choices; the baseline rows reproduce it. The k rows vary the number of preceding trading days that supply control windows. The all-durations row removes the 30-second duration floor for the afternoon open ($k = 5$). The tolerance rows narrow, from 0.5 to 0.1 seconds, the window used to remove the triggering order's own executions from the reference price (Appendix B). The last rows of the continuous-session panels exclude pauses starting between 12:30 and 12:59 PM, whose pre-event window within the afternoon session is short. Standard errors in parentheses and first-stage F-statistics are two-way clustered by firm and date; Anderson–Rubin 95% confidence sets are robust to weak instruments. Treated Events counts the pause events in each estimation sample. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table A.13: Sample Construction for Intraday Pause Events

Stage	Events
Pause Events, One per Stock–Date (All Times of Day)	298,681
Intraday (Continuous Session or 12:30 PM), Resolved Before 2:00 PM	34,164
Duration Filter (Afternoon Open $\geq$ 30 s; Continuous Session $\geq$ 1 s)	17,676
Ex-ante Reference Price Available (Appendix B)	13,780
Daily-Price-Limit Pre-filter	13,772
Estimation Samples (Matched Sets, Controls, Outcome Data)	7,124

Note: This table reports the number of pause events remaining at each step of the sample construction described in Appendix C. For continuous-session events, the reference price is reconstructed from pre-pause transactions using the timestamp rule of Appendix B, with a fallback to the last transaction price on the one-minute grid strictly before the trigger minute. For afternoon-open events, it is the last transaction price on the grid at or before 11:30 AM. Events without a usable same-day reference price (for example, pauses before the first trade of the day) are dropped at that step. The final row is the union of the afternoon-open and continuous-session estimation samples of Table 15.

Table A.14: Robustness Checks for Magnet-Effect Tests

Sample	Observations	Base Rate (bp)	Pause Within Five Minutes (bp)	Log(1 + Traded Value)	Continuation
Full First Section Universe	5,608,714	28.8	−16.38***	−0.028	0.013***
($p \geq 100$ JPY, $\lvert r_5 \rvert \geq 0.8\%$)			(1.60)	(0.038)	(0.001)
Full First Section Universe	6,024,898	26.9	−17.91***	0.141***	0.033***
(No Price Trim)			(1.55)	(0.040)	(0.001)
Full First Section Universe, $z \leq 10$	5,882,321	27.5	−17.00***	0.014	0.024***
(No Price Trim)			(1.47)	(0.038)	(0.002)
$\lvert r_5 \rvert \geq 1.0\%$	2,733,072	49.3	−41.32***	0.024	0.010***
			(2.68)	(0.038)	(0.001)
$\lvert r_5 \rvert \geq 1.5\%$	746,978	122.5	−81.54***	−0.089**	0.002
			(6.55)	(0.042)	(0.002)
Top-500 Stocks	670,684	13.5	−0.49	−0.109***	0.018***
			(3.18)	(0.041)	(0.004)
Outside Top 500	4,938,028	30.9	−18.41***	−0.021	0.012***
			(1.73)	(0.042)	(0.001)

Note: This table reports the coefficient on z from the specifications of Table 16 for alternative samples, defined as in Table 16 unless stated in the row label. Base Rate is the share of stock–minutes in the sample in which a pause starts within the next five minutes. The pause-probability entries (coefficient, standard error, and base rate) are in basis points of probability. In the top-500 group the relevance of z for pause incidence is not established, so its coefficients are not interpreted as responses to anticipation. The top-500 and other-stock regressions each drop their own singleton fixed-effect groups, so their observation counts sum to slightly less than the full-sample count. Standard errors, in parentheses below the coefficients, are two-way clustered by firm and date. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Appendix B. Measuring the Instrument for Continuous-Session Pauses

For a pause that interrupts continuous trading, the instrument requires the price from which the exchange sets the price-limit range, observed before the trigger. During continuous trading, this reference price is the last execution price (Section 2). The difficulty is that the order that triggers the pause may execute in part before the pause begins, so the last recorded transaction may belong to the event itself.

We use identical recorded timestamps to group executions that may belong to the same incoming order. If the final group of transactions before the pause lies within 0.5 seconds of the recorded pause start, we treat it as the triggering order's executions and use the preceding transaction as the reference price. Otherwise, we use the last recorded transaction. Where this reconstruction is unavailable, we use the last transaction price on the one-minute grid strictly before the trigger minute. A tick-precision reference price is available for 2,351 of the 3,438 continuous-session events in the estimation sample (68%). For afternoon-open events, the reference price is the last transaction price on the grid at or before 11:30 AM. Control windows have no trigger: continuous-session controls use the last transaction price on the grid strictly before the matched trigger minute, and afternoon-open controls use the grid price at or before 11:30 AM.

Tightening the tolerance from 0.5 to 0.1 seconds changes the reference price for 80 of the 3,438 continuous-session events (2.3%) and leaves the estimates essentially unchanged: -0.005 (s.e. 0.241) for price-driven pauses and -0.575 (s.e. 0.368) for order-book-evaporation pauses, compared with 0.018 and -0.577 in the baseline (Table A.12).

We also report two descriptive checks, using the lists of up to ten executions that our data record before each pause. We group executions by identical timestamps, omit the first group in each list because its beginning may be truncated, and exclude groups from five seconds before a recorded pause starts through five seconds after it ends; each stock–date–timestamp is counted once. Among 9,475 retained groups with at least two executions, 189 have a high–low range between 1.5 and 2 range steps, and 4 exceed 2 range steps (0.04%; maximum 2.33), with the range step evaluated at the preceding transaction price. This ceiling is what a price-limit range anchored at the preceding transaction implies. Among 121,764 retained groups, 41 changes from the preceding group's last execution to the current group's first execution exceed 2.5 range steps, and none occurs less than 180 seconds after the preceding group. These checks do not uniquely establish the exchange's reference-price rule: the sequential trade quote mechanism (Section 2) may also constrain sweeps, and long intervals between trades do not identify pause-related delays. Nor do they directly validate the timestamp classification or the grid fallback.

Appendix C. Sample Construction for Intraday Pause Events

Table A.13 reports the number of events at each step of the construction of the intraday estimation samples. We start from stock–dates with a single pause and keep pauses that occur during the continuous session or at the 12:30 PM open and end before 2:00 PM, to leave room for the post-event outcome window before the market close. As in Table 3, a pause is classified as an afternoon-open pause if its first record is the afternoon opening call; a pause that begins after the stock has started trading, even within the 9:00 AM or 12:30 PM minute, is a continuous-session pause. We then require pauses to last at least one second during the continuous session and 30 seconds at the afternoon open, and require an ex-ante reference price (Appendix B). This requirement excludes events that occur before the stock's first trade of the day, for which no same-day reference price exists. Finally, we require the range step at the reference price to be strictly smaller than the distances to both daily price limits, and we apply the same pre-filter to control windows.

Control candidates come from the same stock on the ten preceding trading days, excluding stock–dates with any pause. Separately for morning continuous trading, the afternoon open, and afternoon continuous trading, each stock–date is assigned to the candidate event with the smallest trading-day lag. Control reference prices follow Appendix B, and post-event outcomes use the matched event's ending clock time. After requiring usable outcomes, pre-event volatility, and all covariates, we retain

sets containing a treated observation and at least one control. Before estimation, we assign a weight of one to the treated observation and $1/n_c$ to each of its n_c retained controls. Reported event counts are after the removal of singleton fixed-effect groups.

Restricting the intraday analysis to the top-500 sample leaves 1,083 afternoon-open, 6 price-driven, and 50 order-book-evaporation events, with cluster-robust first-stage F-statistics of 5.8, 10.4, and 2.4. The afternoon-open estimate is -0.719 (s.e. 0.807; AR set $[-6.69, 0.51]$), and the AR set for order-book-evaporation pauses is the entire real line. Pauses during continuous trading are therefore studied in the full First Section universe (Section 6.1).

Appendix D. Magnet Effects: Supplementary Evidence

D.1. Pre-trigger dynamics around realized pauses. We use continuous-session pause candidates that survive the event-level filters of Appendix C, together with their available same-stock control windows; this descriptive sample is constructed separately from the IV estimation sample. Windows cover the 30 trading minutes before the trigger minute; windows that would span the lunch break are excluded. For absolute one-minute midpoint log returns and for traded value, we divide the mean over minutes -5 to -1 by the mean over minutes -30 to -11. Each window must have at least ten usable returns in the baseline minutes and three in the final minutes, with positive baseline means. We report unweighted medians across retained windows, without requiring both arms of an event to survive these filters. Before price-driven pauses (514 events and 2,237 control windows), absolute returns and traded value rise to 1.21 and 1.26 times their baseline levels, compared with 0.65 and 0.74 in the control windows. Before order-book-evaporation pauses (793 events and 4,514 control windows), the ratios are 1.04 and 0.96, compared with 0.68 and 0.74 (Wilcoxon rank-sum $p < 0.001$ in each case). The normalized imbalance change, $(\bar{I}_{\text{final}} - \bar{I}_{\text{base}}) / \max\{\overline{|I|}_{\text{base}}, 1\}$, where I is the order-imbalance value of trades signed toward the direction of the pause, is -0.25 before price-driven pauses and -0.16 before order-book-evaporation pauses, compared with approximately zero in the control windows ($p = 0.003$ in both cases).

The control windows are matched on stock and clock time, but not on the preceding price or liquidity shock, and conditioning on realized pauses makes these comparisons descriptive rather than evidence of anticipation. The imbalance change does not establish the sign of order flow in levels, and the rank-sum p-values are not adjusted for dependence across windows.

D.2. Robustness of the magnet-effect tests. Table A.14 re-estimates the tests of Table 16 in alternative samples. The relationship between z and pause incidence is significant in every sample except the top-500 stocks. Among the samples in which z predicts pause incidence, the traded-value coefficient is significantly negative, as a magnet effect implies, only in the sample defined by the 1.5% move threshold, which has 13% of the observations of the baseline sample (-0.089, s.e. 0.042). Without the price trim, the traded-value coefficient is positive (0.141); excluding observations with $z > 10$ reduces it to an insignificant 0.014. All continuation point estimates are positive, and none is significantly negative in the direction predicted by a magnet effect. In the top-500 sample, the relationship between z and pause incidence is not statistically distinguishable from zero, so we do not interpret its trading coefficients as evidence about anticipation.