Do Trading Pauses Protect Retail Investors? *

Don Noh†

April 2023

Abstract

This paper studies retail investor trading behavior around stock-level trading pauses in Korea, which are triggered after extreme intraday price movements. The stated goals of this market intervention tool are investor protection and price stability. However, I find no significant improvement in terms of 5-minute realized volatility, daily high-low price range, and frequency of extreme price movements. Instead, I find unintended negative consequences on retail investors. Because trading pauses are salient events, retail investors trade excessively once pauses are triggered, and they trade in the wrong direction by betting on price continuation. Trading pauses become a window of wealth transfer from retail to institutional investors, and estimates for the wealth transfer and trading costs together amount to approximately USD 350 million per annum.

* This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant. I am grateful to my advisors Motohiro Yogo, Wei Xiong, and Markus Brunnermeier for their guidance and support. I also thank Stella Park, Sangmin Oh, Jihong Song, Francisco Cabezon, Jonathan Payne, Simon Schmickler, Nicolas Hommel, Ziang Li, Philippe van der Beck, Donghwa Shin, Batchimeg Sambalaibat, Caio Almeida, Moritz Lenel, Yacine Aït-Sahalia, Atif Mian, Jiwon Lee, Sohye Yoon, Amora Elsaify, and participants of the Princeton Finance Workshops for feedback.

† Princeton University (dnoh@princeton.edu).

1 Introduction

On March 9, 2020, when COVID-19 induced fears overtook global investors, circuit breakers at the US stock exchanges kicked in. There were three more days in that month during which such market-wide circuit breakers were triggered. Above and beyond the news that the S&P 500 Index fell by more than 7% within a single trading day—the threshold for a level 1 circuit breaker—the term *circuit breaker* conjures up an ominous emotion, and certainly makes for a good headline. Given the daily volatility of stocks, the economic news contents in the headlines (1) *S&P 500 Index fell by 6.5%* and (2) *S&P 500 Index triggered a level 1 circuit breaker* must not be substantially different. However, the salience of this regulatory measure makes the latter caption simply harder to ignore.

Both market-wide circuit breakers and stock-level trading pauses are triggered by either extremely positive or negative intraday returns, and are meant to serve as a speed bump to panic selling and manic buying. They aim to achieve this by providing market participants time to reassess the situation and react, thereby protecting investors from extreme price swings.[1] In this sense, salience is a feature rather than a bug. But what types of reactions are these interventions soliciting? What if they stir up more animal spirit than liquidity provision? Should we be worried about any unintended consequences?

In this paper, I exploit high frequency, investor type-level trade flow data from the Korean stock market to show that one such unintended consequence is a consistent and systematic wealth transfer from retail investors to institutional investors. I come to this conclusion by documenting that (1) retail investors as a whole make bets on price continuation around trading pauses—net buying (selling) stocks that trigger a pause by appreciating (depreciating) in price—and that (2) prices revert after these trading pauses. Effectively, retail investors are net buyers at the peak around an upward trading pause and net sellers at the trough around a downward trading pause. On the other hand, foreign institutions absorb more than 80% of these net flows and emerge as the biggest winners through trades around trading pauses. This pattern is strikingly consistent across tens of thousands of trading pause events.

As a benchmark, I compute the transfer of wealth from retail investors to foreign institutions under the assumption that wealth is marked to market at the end of the trading day using the daily closing price of a stock. For instance, if retail investors net bought one share at \$100 around the pause and later this share closes at \$90, this would imply a daily loss of \$10 under this simplifying assumption. Back-of-the-envelope calculation suggests that in the first two years of rule adoption, retail investors as a whole transferred around USD 150 million of wealth to foreign institutions during stock-days on which trading pauses were triggered. Moreover, retail investors' heightened trading activity around a trading pause is also subject to brokerage fees

[1] Deutsche Börse describes their trading pause rules as "protective mechanisms" and Korea Exchange states that trading pause rules are put in place to "prevent damages to investors."

and trade taxes that amount to 0.3% of total trading value.[2] These costs amount to an additional USD 550 million, making the estimated losses exceed USD 700 million. Indeed, these are not causal estimates because it may well be the case that retail losses could be similar on extreme return days—days during which prices move outside the $\pm 10\%$ band relative to opening prices—even without the trading pause rule. However, even after using such pre-rule extreme return days as reference points, the resulting estimates do not change significantly.

To be more specific about the setting, I study a stock-level, 2-minute trading pause rule—called the volatility interruptions (VI) mechanism—that was introduced to the Korean stock market on June 15, 2015. Abstracting from a few details, VI is triggered when a security listed on the Korea Exchange (KRX) moves more than 10%, either up or down, relative to its opening price.[3] Once VI is triggered, continuous trading halts for 2 minutes and an auction phase begins. These 2-minute VI auctions work in the same way as do opening and closing auctions in NYSE or NASDAQ: traders can add, modify, or cancel orders, these orders are accumulated to the order book, and they are cleared only at the end of the auction phase. The state of the aggregate order book is public so that anyone can observe the indicative clearing price and volume throughout the auction phase. As with opening and closing auctions, the goal is to concentrate liquidity and to promote price discovery.

The KRX data prove to be particularly well suited for several reasons. The first is that KRX provides exact identifiers for investor types, such as retail investors and foreign institutions, at the *individual trade level*. For every single executed order, I observe the investor types of the buyer and seller, on top of the usual variables like order type, order volume, and executed price.[4] The second is that the investor type-level trade flows are exhaustive. Earlier papers such as Barber and Odean (2008) and Hirshleifer et al. (2008) use retail flows from brokerage firms. More recent innovations by Boehmer et al. (2021) also enable researchers to impute retail flows from the US trade and quote (TAQ) data. However, neither of these methodologies capture the exhaustive flows of different investor types. Because KRX is the sole stock exchange on which all stock trades happen, I can make definitive statements about the aggregate flows between investor types. Lastly, the KRX data exist both for the periods before the rule adoption (pre-rule period) and after the rule adoption (post-rule period). This allows for the sampling of *pseudo-pauses*—a stock's breaching of the pause trigger limit, $\pm 10\%$, during the pre-rule period—which can serve as control events.[5]

With this detailed dataset, I first document that retail investors' trading volume increases disproportionately during the minutes immediately following the trading pauses. Earlier studies (Christie et al. 2002; Hautsch and Horvath 2019) have also found heightened volume following

[2] This is a conservative estimate using a relatively low brokerage trading fee. Losses are translated to USD terms using the approximate average KRW/USD exchange rate of 1,100 during the years 2014-2016.

[3] The exact rule is discussed in more detail in Section 2.1

[4] While this is true in principle, I have trade-level data from June 2014 and May 2016. This is due to funding constraints. For periods extending to May 2018, I use 1-minute frequency data instead.

[5] The term and concept are adopted from Lee et al. (1994) and Hautsch and Horvath (2019).

trading halts or pauses. I refine these findings by showing that more than 95% of this extra activity is coming from retail investors. Similar to an exercise in Seasholes and Wu (2007), which shows that attention-grabbing effects of halts are weaker when there are multiple contemporaneous halts, I demonstrate that the extra retail trading garnered by a trading pause decreases monotonically with the number of preceding pauses on the same day. This adds weight to the narrative that retail investors are reacting to the salience of this regulatory device.

In addition to these ex-post patterns, I also present notable changes in retail investors' ex-ante trading behavior near the 10% VI thresholds. Generally speaking, retail investors act as liquidity providers in the Korean market: net selling stocks whose prices have gone up and vice versa for stocks that depreciated.[6] This pattern holds regardless of whether I use contemporaneous—with respect to retail flows—returns or lagged returns, and whether I look at shorter horizons (hours) or longer horizons (days and weeks). This is unsurprising given that institutional flows are informed and persistent. During the pre-rule period, retail investors' tendency to net sell becomes monotonically stronger as a stock appreciates relative to its opening price. However, in the post-rule period, this monotone pattern breaks down and retail investors suddenly become net buyers when a stock's price nears the $+10\%$ threshold—i.e., when a stock appreciates by 7, 8, or 9% relative to its opening price. Why do retail investors suddenly become buyers of such *expensive* stocks from being sellers? I attribute this to the anticipation of the ex-post trading pattern and overconfidence. If a trader expects to correctly foresee an upward VI and the ensuing late-arriving net buy orders, this trader has an incentive to buy. By backward induction, this creates a cascade of incentives to net buy when the probability of upward VI rises. I demonstrate the existence of this incentive by constructing a trading strategy that successfully implements this idea.

Notwithstanding these concerning trading patterns, circuit breaker-like rules are ubiquitous. This leads us to expect that trading pause rules introduce other benefits to overall market conditions that are not yet discussed. I exploit the arbitrariness of the pause-triggering threshold to identify the rule's effects on volatility and liquidity. The arbitrarily chosen 10% threshold is *tight* for a security with a high daily volatility whereas it is practically nonexistent for a security like Samsung Electronics' common stock. In the benchmark specification, stocks are sorted by return volatility in the past 60 trading days. Then, the top quintile is considered the treatment group while the bottom quintile is considered the control group. The outcome variables used are monthly averages of 5-minute realized volatility, high-low price range, bid-ask spread, order book depth, Amihud measure, and the number of $\pm10\%$ breaches during a given month. I find statistically insignificant effects on measures of volatility and mixed effects on measures of liquidity. For robustness, I also sort by the number of extreme returns and daily high-low price range in the past 60 trading days, and find similar results. One concern is that more volatile stocks may be more sensitive to overall market conditions relative to less volatile stocks. To

[6] Barrot et al. (2016) find the same pattern in the French stock market.

mitigate this concern, the main specification includes quintile-specific slopes to market volatility and market returns as controls.

Altogether, the results demonstrate how the interplay of behavioral biases (salience and extrapolation) and regulatory factors (trading pause rule) confounds this seemingly simple exchange rule. As epitomized by the Lucas critique (Lucas 1976), economists heed much attention to endogenous responses of market participants. I provide an example where failing to account for *behavioral* responses to rule changes attenuates, and even reverses, the intended effects.

My results are closely related to the main narrative of Seasholes and Wu (2007) and Chen et al. (2019b): a group of naive traders, who are often small retail investors, flock to stocks that recently breaches an upward limit, and sophisticated traders who correctly anticipate this actively buy near the upper price limits only to unload positions to the latecomers. I further the understanding about what drives both the ex-post and ex-ante incentives of retail trading around pauses. In addition, the high resolution investor type-level data reveal that a significant amount of wealth transfer happens from retail investors to institutions, and also allow me to quantify this as well.

This paper also contributes to the literature on attention-induced trading by individual investors and its perils (Barber and Odean 2008; Engelberg et al. 2012; Barber et al. 2022). Most papers in this area use performance measures related to hypothetical returns to measure trading losses. Because exhaustive retail flows are observed at a high frequency, this paper is able to compute the amount of wealth outflow directly. Similar to the findings in Barber et al. (2009), retail investors' performance suffer due to unnecessary trading, while foreign institutions reap most of the direct benefits in my setting as well. One takeaway from Barber et al. (2022) is that how trading venues are designed and presented to retail investors matter for the trading outcome of retail investors. Results here echo this lesson.

Various other papers also study the intended and unintended effects of trading pauses. Studies on trading pause rules have faced challenges to identifying the effect of adopting these rules because exchanges around the world often introduce them in a non-staggered manner—i.e., to all stocks on the same day. A common workaround has been looking at volatility and liquidity around realized pauses. Studies have found that markets experience heightened activity around pauses, contrary to the exchanges' aim of cooling-off markets (Lee et al. 1994; Christie et al. 2002; Hautsch and Horvath 2019). However, this alternative approach faces the limitation that only *local* statements can be made. Markets may be more turbulent around realized pauses, but they may experience less pause-triggering price movements in the first place. I show that this is unlikely to be case through a new workaround that exploits the arbitrariness of the pause-triggering threshold.

Studies have also considered how the endogenous reaction of traders to circuit breaker-like rules can affect trading during normal times. A classic example is the "magnet effect" which refers to a situation where traders drive the price towards the limit when prices are close to the interruption thresholds (Subrahmanyam 1994; Cho et al. 2003; Yan Du et al. 2009). On downwards swings, for instance, leveraged buyers may unravel positions in expectation of inability to rebalance their portfolios, thereby adding to the selling pressure (Chen et al. 2019a). This paper's result on the ex-ante trading behavior of retail investors suggests that they indeed behave in a way consistent with the magnet effect hypothesis.

2 The Korean Stock Market

2.1 Notable Aspects

Korea Exchange (KRX) is the sole exchange through which all stocks, ETFs, and derivatives clear in Korea. It consists of three divisions: the KOSPI market division, the KOSDAQ market division, and the derivatives market division. KOSPI houses ETFs and larger blue-chip stocks such as those of Samsung Electronics, Hyundai Motor, POSCO and LG Electronics, while KOSDAQ houses smaller technology stocks. There are 793 companies listed on KOSPI and 1,497 companies listed on KOSDAQ.[7]

The total market capitalization of stocks listed on KRX amounts to 1.67 trillion USD. This places Korea's public equity market in between those of Germany and Australia in terms of size. More interestingly, the average monthly trading volume is 277 trillion USD, trailing only the US, China, and Japan in this department. On top of its considerable size and activity, the Korean stock market possesses several institutional details that make it an attractive setting to study retail investor behavior.

Active Retail Participation. First, the Korean stock market is a retail-driven market. As shown in Figure 1, the retail share of monthly trading volume in KRW terms has hovered around 65% and even reached highs of 80% during the pandemic-induced retail trading surge. The same statistic in the US moved from around 15% to 20% during the same periods.[8] According to the Korean Securities Depository (KSD), a custodial service provider, retail investors account for 28% of total stock holdings. Consistent with Odean (1999), Korean retail investors trade much more intensively compared to institutional investors.

[7] All market statistics such as number of listed companies, market capitalization, or trading volume are as of September 2022.

[8] Source: Bloomberg Intelligence.

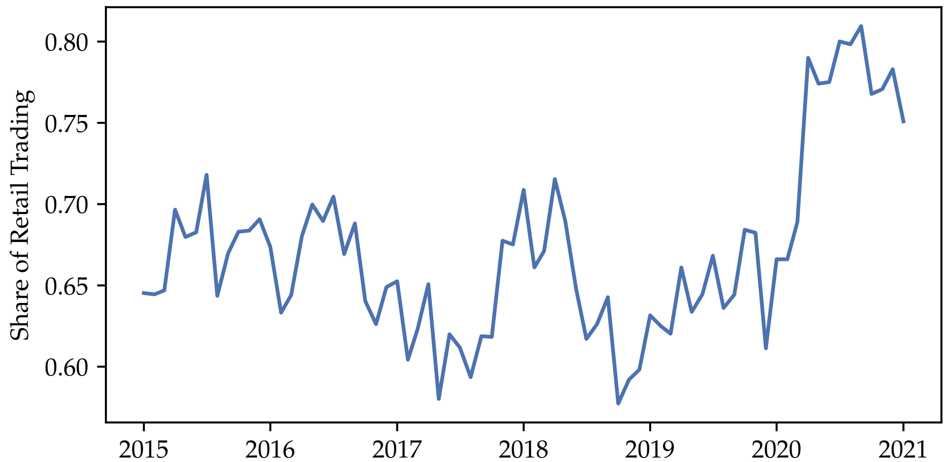

Figure 1. Retail Trading Volume Share

This figure plots the share of retail investors' daily trading volume relative to total trading volume in the Korean stock market. Both total trading volume and retail trading volume are measured in KRW terms and aggregated at the monthly frequency. The ratio of the sums are reported in the figure.

Known locally as "ants," Korean retail investors are both an object of media's ridicule and market's scrutiny, respectively due to their poor performance and sizable influence. As Kim and Kim (2022) finds, by analyzing approximately 20,000 individual brokerage accounts, Korean retail investors indeed demonstrate many of the well documented behavioral traits such as overconfidence (Scheinkman and Xiong 2003), disposition effect (Shefrin and Statman 1985), penchant for lottery-like stocks (Bali et al. 2011), and herding. Thus, interesting retail trading patterns come in both variety and quantity in the Korean stock market.

However, not all retail investors should be considered equal. The largest 0.5% of retail accounts, whose balances are larger than 1 billion KRW, constitute around half of total holdings amount. Similar to Chen et al. (2019b), Kim and Kim (2022) also find that larger accounts are less prone to the aforesaid behavioral biases. Given these facts, it is more reasonable to consider the aggregate retail trade flows as coming from a mix of *sophisticated retail investors* and *naive retail investors*.

Investor Type-Level Trade Flows. The most attractive part of the KRX data is that KRX provides identifiers for the investor type at the *individual trade-level* through their trades and quotes (TAQ) data, albeit at a significant price. The provided investor types are: (1) retail investors, (2) foreign institutions, (3) proprietary traders, (4) asset managers, (5) pension funds, (6) banks, (7) insurance companies, and (8) private equity. All investor types except retail investors and foreign institutions refer to domestic institution types.

At a more reasonable price, KRX also provides trade flows aggregated at the investor type level at 1-minute frequencies. This is the data that most of the analyses in this paper will rely on. For example, I can observe how many shares of Samsung Electronics retail investors bought and sold as a group between 9:01 and 9:02 a.m. on March 15th, 2020. This allows me to compute, for instance, retail net buys (shares bought minus shares sold) or retail volume during any 1-minute interval within my sample period. The same applies to all other provided investor types.

This establishes two important advantages over the previously used data on retail flows: exhaustiveness and precision. Various studies using brokerage data (Hirshleifer et al. 2008, Luo et al. 2022) are subject to the constraint that they have to work with a representative sample of retail trades. In a recent study, Boehmer et al. (2021) provides an innovative way to identify retail orders from the US TAQ data using regulatory restrictions on price improvement. However, the method is still a conservative imputation that applies to marketable retail orders, rather than marketable and limit orders together. With the KRX data, I can safely assume that I am working with *the* retail flow.

Using such data, I can immediately document the following interesting facts, for instance: Figure 2a shows that retail investors were net buyers throughout the pandemic-induced trading frenzy and Figure 2b shows that they are net buyers in the morning.

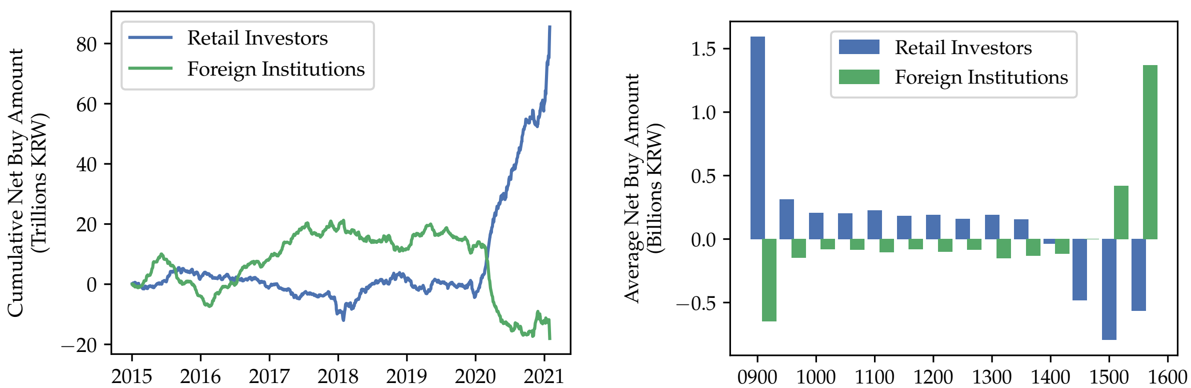

(a) Cumulative Retail Net Buy

(b) Retail Net Buy by Time of Day

Figure 2. Retail Net Buy Patterns.

Panel (a) plots the cumulative net buy of retail investors (in green) and foreign institutions (in blue) from Jan 2015 to Mar 2021. Retail (foreign) net buy is measured each day by summing up retail buy volume in all stocks in KRW terms and subtracting the sell volume in all stocks. Then, the cumulative sum of these daily net buys are plotted. Panel (b) plots the average net buy volume during a given 30-mintue interval over all stock-day observations. Net buy is defined analogously, but aggregated at each stock-day-time intervals.

Volatility Interruption (VI) Mechanism. VI is a trading pause rule that the KRX adopted in June 15th, 2015. It is triggered when the to-be-executed trade price lies outside the "pre-defined price range," which will be discussed in the next paragraph. Once a VI is triggered, continuous trading is paused for 2 minutes and the "auction phase" begins. During the auction phase, orders can be entered, changed, and cancelled, but will not be immediately executed. Each of these actions will be accumulated on the limit order book and will affect the prevailing supply and demand curves. At any moment during the auction phase, indicative price and volume—the intersection of supply and demand given the state of the order book at that moment—is broadcasted to market participants. At the end of the 2-minute auction phase, the stock emerges with a single equilibrium price and continues to trade normally. VI differs from a trading *halt* in this respect: a trading halt precludes trading for the rest of the trading day.[9]

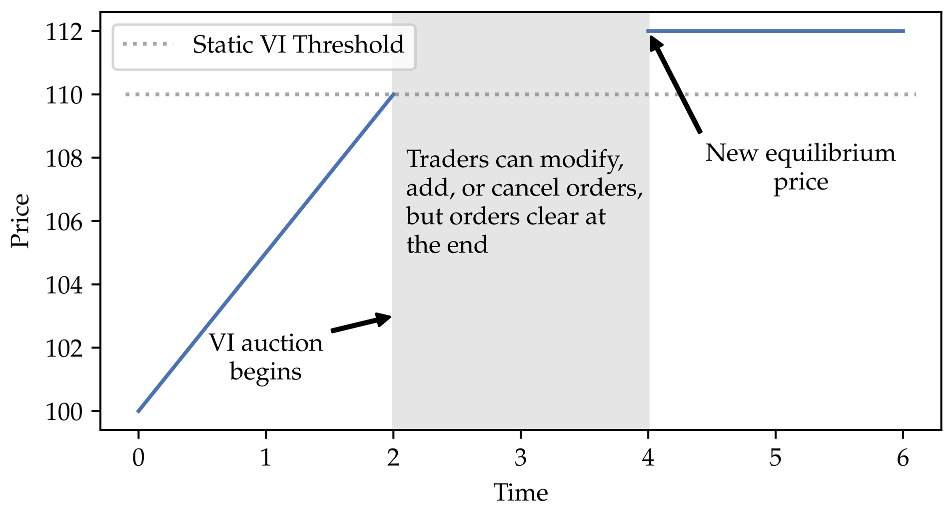

Figure 3. Description of VI Auction Mechanism

This figure illustrates what happens when a static VI is triggered. The blue line is an example path of a stock's price while the gray dotted line is the static VI threshold (10% relative to opening price). Once the price breaches this threshold, a 2-minute VI auction phase begins. At the end of the 2 minutes, the uniform price that clears the market becomes the new equilibrium price and continuous trading resumes. This pictorial description is adopted from Deutsche Börse's website and modified for my use.

There are two types of VIs at KRX: static and dynamic. A static VI is triggered when the to-be-executed trade price lies outside ±10% relative to the reference price which is defined as "the most recent auction price." This will be the price determined by the most recent of closing auction, opening auction, or VI auction. For most situations, it is safe to assume that the reference price is the opening price, unless a stock has already triggered a VI on the same day. A dynamic VI is triggered when the to-be-executed trade price lies outside ±3% relative to its previously executed price for constituents of the KOSPI 200 and ±6% for other securities. It is triggered only if a *single order* moves the potential trade price outside the price range. Thus, a

[9] The description of the VI mechanism and Figures 3 and 4 rely heavily on the Deutsche Börse's descriptions of the same rule.

static VI is targeted toward large cumulative intraday price moves while a dynamic VI targets situations that resemble fat-finger trades. Going forward, I refer to a static VI triggered by an upward price movement as a static-up VI. Other three cases are abbreviated similarly.

KRX adopted the dynamic VI rule on September 1, 2014 and the static VI rule on June 15, 2015. They were introduced as a way to mitigate the impact of widening the daily stock-level price limits, thresholds at which stocks halt trading for the entire day, from ±15% to ±30%. Since June 15, 2015, all stocks and ETFs are subject to daily price limits of ±30%, dynamic VI, and static VI rules. All rule changes applied to all listed stocks from the adoption date onwards, which poses challenges to identification.

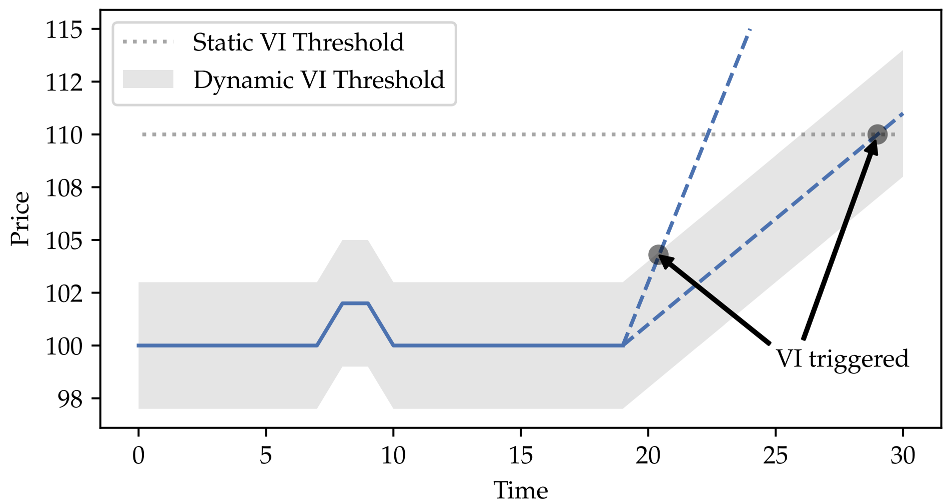

Figure 4. Description of Static VI vs. Dynamic VI

This figure illustrates when static and dynamic VIs are triggered. The blue line is an example path of a stock's price, the gray area is the dynamic VI threshold corridor (3% relative to the latest execution price), and the gray dotted line is the static VI threshold (10% relative to the opening price). The two dotted blue lines represent two possibilities: (1) price breaches the dynamic corridor by a *single trade* (left dot) or (2) price breaches the static threshold anytime during the trading day (right dot). In either case, a VI auction as described in Figure 3 ensues. This pictorial description is adopted from Deutsche Börse's website and modified for my use.

According to the KRX website, the stated goal of the VI rule is to "prevent damages to investors due to excessive price movements." Similarly, the Deutsche Börse classifies the VI rule as a "protective mechanism." These statements do not explicitly mention retail investors. However, the emphasis on cooling-off of sentiment and providing time to react suggest that non-sophisticated investors were likely targeted at the conception of VI-like rules.

2.2 Korea Exchange (KRX) Data

As of September 2022, KRX data covers 2,411 common and preferred stocks, 622 ETFs and various other securities listed on KOSDAQ and KOSPI markets. All data listed below can either be accessed directly on KRX's website or can be purchased from KRX, and span at least one year of the pre-rule period.

1-Minute Frequency Trade Flow and Price Data. 1-minute frequency data that I work with runs from June 2014 to May 2018. This data can be considered redundant because they can be computed from trade-level data. I rely on the 1-minute frequency data for longer run analysis because trade-level data cost up to ten times the 1-minute frequency data. Trade flow data includes buy/sell volume and buy/sell value (KRW) at 1-minute frequency by investor type. For any 1-minute interval, I can compute the average price retail investors paid for Samsung Electronics common stock by dividing retail investors' buy value by their buy volume. Price related data include open, close, high, low prices, traded volume/value, liquidity measures, and order imbalance measures. All available variables are listed in Table A1.

Trade-Level Data. The trade-level data that I work with runs from June 2014 to May 2016. For every trade that is executed on KRX, I observe the investor type of the buyer and seller, volume and price of buy/sell orders, order types (market, limit, stop limit, etc.), order time (microseconds), and snapshots of the order book (up to 10 levels) at trade execution moments. Notice that this is different from the quote-level data which records all changes to the order book. Other variables are listed in Table A1.

Other Firm-Level Data. KRX also provides various daily statistics such as short volume, short interest, book-to-market ratio, dividend yield, and shares outstanding on its website. For regulatory reasons, foreign institutional holdings are also reported at a daily frequency, but holdings of other investor types are not directly reported. Another source of firm-level fundamentals is the Data Analysis, Retrieval and Transfer System (DART) which is the EDGAR counterpart in Korea.[10] Other variables are listed in Table A1.

Summary Statistics. I begin with some relevant statistics that give us a better perspective. Figure 5 plots the number of static and dynamic VI occurrences per day. As one may expect, the number of VIs spike during turbulent days. During the COVID-19 crash the number of VIs exceed the number of listed stocks because a stock may hit multiple VIs in a single day. Because a 10% intraday price movement is large, one may expect that VIs are very rare events. As we see in Table 1, the average daily volatility for liquid stocks is around 3%. Using an inflated number of 5% suggests that a 10% intraday move should be two standard deviations under the

[10]DART provides excellent automation and API access through OpenDART.

assumption of normality. The median daily number of static VIs is 104—around 5% relative to the number of listings—which is in the ballpark of our inflated estimate, but it is true that the VI occurrences are more frequent than what a normal return distribution would suggest. This high median suggests that VIs are quite frequent during normal times as well and gives us ample statistical power.

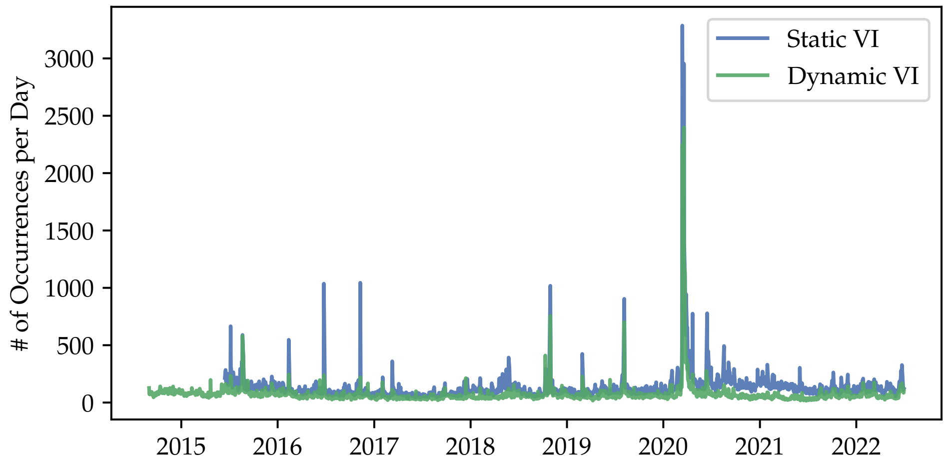

Figure 5. Number of Daily VI Occurrences by VI Type

This figure plots the number of daily VI occurrences by VI type. Dynamic VIs were introduced in September 2014 and static VIs in June 2015. Repeated VIs during the same day are counted separately for this figure.

As discussed at the beginning of this section, there are more than 2,000 listings on KRX, which is a large number relative to total market capitalization. This means that most of the stocks will be small and often illiquid. For this reason, I use a subset of 980 stocks (1) that traded every day from 2014 to 2018, (2) whose average market capitalization larger than 50 KRW billion, and (3) whose average daily trading value is larger than 0.5 KRW billion. Table 1 provides a summary of the stock-month level observations.

Table 1. Summary Statistics of Listed Stocks

	Count	Mean	SD	1%	25%	50%	75%	99%
Market Cap. (KRW trillion)	84,585	1.55	10.85	0.03	0.10	0.20	0.64	21.70
Avg Daily Return Vol. (%)	84,585	2.99	1.33	1.12	2.04	2.67	3.62	7.26
Market Beta	84,585	0.93	0.68	−0.86	0.52	0.92	1.34	2.63
Price-Earnings Ratio	84,581	48.59	586.19	0.00	0.00	12.06	27.38	527.10
Price-Book Ratio	84,585	2.25	3.85	0.28	0.85	1.40	2.46	14.19
Foreign Holdings (%)	84,585	9.63	12.78	0.00	1.67	4.22	12.56	60.30

This table reports the summary statistics for the subset of chosen stocks. Each observation is at the stock-month level. Market capitalization, share of foreign institutional holdings, P/E ratio, and P/B ratios are averaged within the same stock-month. Daily return volatility is the standard deviation of daily returns of a given stock during the month of interest. Market beta is the rolling 60 day market beta of a given stock as of the first day of the month of interest.

Definitions. *Net buy* by investor type is a quantity that will appear repeatedly. It is defined as the number of shares purchased by an investor type during an interval minus the number of shares sold by the same investor type. In order to make the quantity comparable across different stocks, net buy is often normalized by shares outstanding. In few of the analyses, in Figure 2a for example, I use KRW value instead of number of shares.

$$\mathrm{NB}_{ijt} = \#\text{ shares of stock } i \text{ bought minus sold by investor type } j \text{ during interval } t$$

$$\mathrm{nb}_{ijt} = \frac{\mathrm{NB}_{ijt}}{\#\text{ shares outstanding}_{it}}$$

Because markets have to clear during any interval $t \to t+1$, we must have

$$\sum_j \mathrm{NB}_{ijt} = 0 \text{ and } \sum_j \mathrm{nb}_{ijt} = 0 \text{ for any stock } i \text{ and time } t.$$

For this reason, unconditionally we should expect to see zero net buys for any randomly chosen (i, j, t). This means that it is more often the case that negative retail net buy, which mechanically implies positive institutional net buy, over a period is associated with rising prices. Another related flow definition is *aggressive buy*, which is defined analogously as:

$$\mathrm{AB}_{ijt} = \#\text{ shares of stock } i \text{ bought by investor type } j \text{ during interval } t \text{ at } \textit{ask price}$$

$$\mathrm{ab}_{ijt} = \frac{\mathrm{AB}_{ijt}}{\#\text{ shares outstanding}_{it}}$$

By definition, an aggressive buy cannot move the executed price downward. Of course, if bid and ask quotes adjust downwards as time goes by, even positive aggressive buy can be associated with negative price movements. However, this situation is rarer. It should be noted that aggressive buy includes both limit buy orders that cross the spread and market buy orders.

In the following sections, the majority of the results pertain to static VIs whose reference price is the opening price of the stock in question. For this reason, it is convenient to define *return from open* as the following:

$$\text{return from open}_{it} = \frac{P_{it}}{P^o_{it}} - 1$$

where P^o_{it} is the opening price of stock i on the day in which time t is included. This means that the first static VI occurs when return from open reaches $\pm 10\%$ for the first time.

3 Ex-Post Effect of VI on Retail Trading

The first thing I document is that VIs are salient events and they induce retail investors to trade more intensively. I also document that retail investors bet on price continuation: they net buy on upward VIs and net sell on downward VIs. I start by looking at the retail trading volume and retail net buy around static VIs.

3.1 Stylized Facts

A stock that has risen or fallen by 10% relative to its opening price faces a selection problem. Such a big move may be driven by a big news event, liquidity shock, or sudden change in sentiment. Because I have data from periods before the VI rule was put in place, I use the observations from those periods as reference points. Similar to Lee et al. (1994) and Hautsch and Horvath (2019), I call these pre-rule threshold breaches *pseudo-pauses*. For periods before rule adoption, from June 2014 to May 2015, I sample the minutes around the first breach of $\pm 10\%$ for each stock. For periods after rule adoption, from July 2015 to December 2016, I sample the minutes around the first static VI.

Because both volume and net buys at 1-minute intervals are very noisy, I aggregate them at 10-minute intervals instead. One thing to note is that the interval from minute 0 to minute 10 starts with trades that are executed *after* the VI. Recall that once a VI is triggered, a 2-minute auction phase begins. This means that the first 10-minute interval completely includes the auction phase, and none of the trades that happen before the pause is triggered. If abnormal trading and net buying by retail investors happen during the auction period, they will be reflected in the interval from minute 0 to minute 10. I present the results for retail volume in Figure 6 and for net buy in Figure 7.

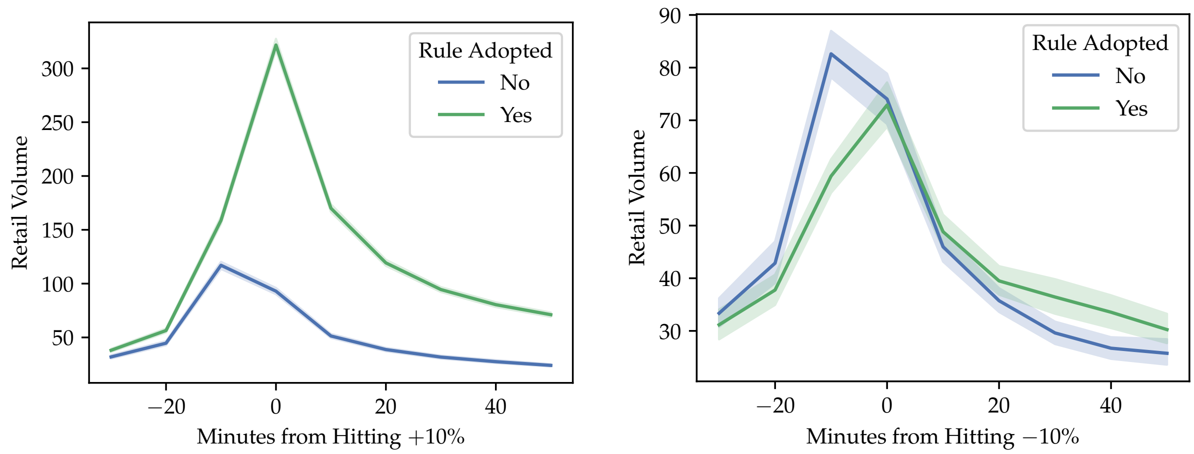

(a) Static-Up VI

(b) Static-Down VI

Figure 6. Retail Volume around Static VIs

This figure plots the average retail volume during 10-minute bins around VIs and pseudo-pauses. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, volume is expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). Panel (a) plots retail volume around static-up VIs. Panel (b) plots quantities around static-down VIs. 95% confidence intervals are constructed by bootstrapping.

Figure 6 shows spikes in volume in the first 10-minute interval that comes immediately after the breaching of 10% thresholds. We are interested in the height of the green line relative to that of the blue line at $x = 0$. Compared to a pre-rule 10% breach, retail volume is around 200 bps of shares outstanding higher. This is a large magnitude considering the brevity of the interval. Furthermore, activity stays heightened for the next hour or so.

The first thing to note is the difference between the upward breach and the downward breach. Retail investors appear to react much more to a static-up VI compared to a static-down VI. This may be due to the fact that there is a large regulatory hurdle for retail investors to short. Retail investors can always make attention-induced buys, but cannot make attention-induced sells without owning the stock.

The second interesting pattern is that volume peaks immediately *before* the breaching of 10% when VI was not yet introduced, and it peaks immediately *after* the breach once VI is in place. A threshold breach happens when an extreme price movement happens. This means a sudden large activity is likely to precede the threshold breaches. Accordingly, we see high volume at $x = -10$ for all cases (up-down and pre-post). The main difference in Figure 6a is that volume is pushed up even further following a static-up VI, while it recedes after the 10% breach in a pseudo-pause. Concerns about reverse causality will be discussed further later in the section.

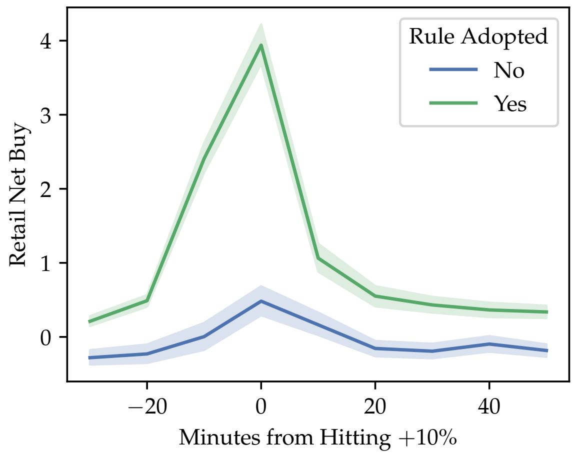

(a) Static-Up VI

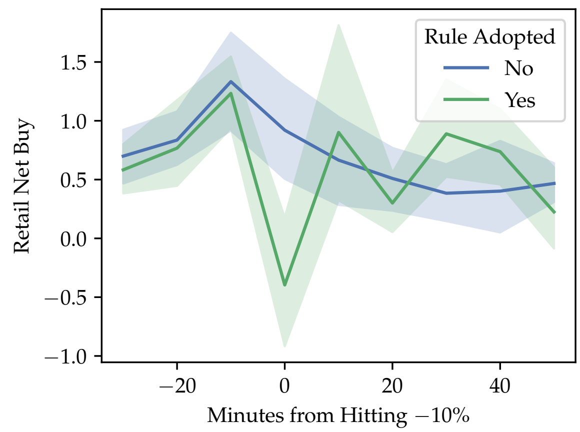

(b) Static-Down VI

Figure 7. Retail Net Buy around Static VIs

This figure plots the average net buy during 10-minute bins around VIs and pseudo-pauses. Retail net buy is computed as retail buy volume minus retail sell volume. Net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, net buy is expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). Panel (a) plots retail volume around static-up VIs. Panel (b) plots quantities around static-down VIs. 95% confidence intervals are constructed by bootstrapping.

Figure 7 repeats the same exercise with retail net buys. Again the pattern is cleaner with upward breaches. Panel (a) shows that that retail investors are the net buyers of a stock around the breaching of 10%. In other words, more retail investors are choosing to buy a stock that has significantly risen in price from the opening of the market. While the absolute magnitude of the spike is small at around 2 bps of total shares outstanding, this is *net* retail buy during a 10-minute interval; unconditionally, this number should be zero.

While clear patterns appear in this unconditional averaging exercise, it is necessary to account for temporal patterns and other stock-level controls. We will see in the following section that adding controls do not affect the pattern.

3.2 Empirical Results

Panel OLS. I start by using a simple panel ordinary least squares (OLS) regression to confirm the above patterns. For each stock, I first split the time series into 10-minute bins so that each observation (i, t) is stock i in 10-minute bin t. To clarify, t is not the time of the day, but rather the t-th 10-minute interval during the entire time-series. For days during which a static-up VI occurs, the 10-minute bins are constructed so that the moment of VI occurrence lies at the beginning of a 10-minute bin. Figure 8 illustrates an example of this binning. If VI occurs at 9:36

a.m., trade flows following the VI will be included in the interval 9:36 → 9:46 a.m. and flows immediately preceding the VI will be included in the interval 9:26 → 9:36 a.m. The minutes preceding 9:06 a.m. are discarded because 9:00 a.m. is the opening time.

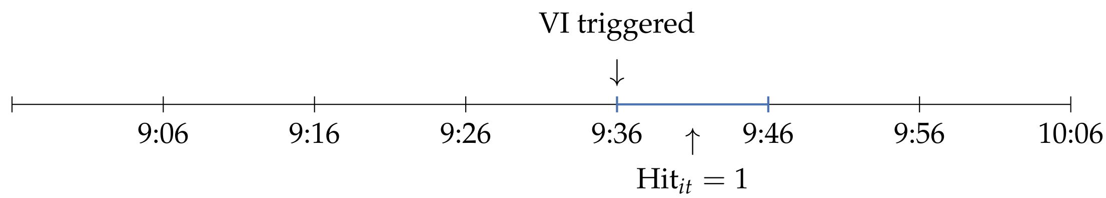

Figure 8. Interval Construction

This figure illustrates how 10-minute intervals are constructed from the trade flow data. The figure is an example in which ±10% breach happens at 9:36 a.m. Volume and net buy are aggregated at 10-minute intervals around this breach. Later, the first interval following the breach is assigned $\text{Hit}_{it} = 1$ and other intervals are assigned $\text{Hit}_{it} = 0$.

A 10% threshold breach in the pre-rule period, a pseudo-pause, is treated in the same manner. Once this bin assignment is done, the following dummy variables are assigned:

$$\text{Hit}_{it} = 1 \text{ if stock } i \text{ breached +10\% at } t$$
$$\text{Post}_{it} = 1 \text{ if interval } t \text{ is part of the post-rule period.}$$

In a panel regression, the coefficient on Hit_{it} captures the effect of breaching the 10% threshold, while the coefficient on $\text{Hit}_{it} \times \text{Post}_{it}$ captures the effect of experiencing a VI in addition to the effect of a threshold return breach. The panel OLS specification is as follows:

$$y_{it} = \gamma_1 \cdot \text{Hit}_{it} + \gamma_2 \cdot \text{Post}_{it} + \beta \cdot \text{Hit}_{it} \times \text{Post}_{it} + \Gamma' \mathbf{X}_{it} + \text{Hour FE} + \varepsilon_{it}. \quad (1)$$

The outcome variables used are total volume, retail volume, foreign institution volume, retail net buy, and foreign institution net buy of stock i during an interval t. The vector $\mathbf{X}_{it}$ includes control variables such as log market capitalization, average market beta in the past 60 trading days, average Amihud measure in the past 60 trading days, return in the past 20 days, and return in the past 200 days. Lagged flows are also included as controls when the dependent variables are flows because they tend to be persistent.

The coefficient of interest is β. Columns (1), (2), and (3) of Table 2 show that trading volume is significantly higher during the first 10-minute after following a static-up VI even relative to a pre-rule pseudo-pause. What is also notable is that almost all the effect is accounted for by rise in retail volume. The magnitude 183.9 bps for retail volume lines up well with the results of the previous subsection.

Columns (4) and (5) show that retail investors are net buyers during this first interval and the foreign institutions are handling most of this net buys. What is also interesting is that foreign institutions' rise in volume and rise in magnitude of net buy are almost the same: 1.14 bps vs. 0.9 bps. This means that foreign institutions are determined in their trade direction, and are mostly selling without buying. The magnitudes for both net buy regressions are small because the flows are aggregated at the investor type level.

Table 2. Panel Regression Coefficient Estimates

	Total Vlm	Retail Vlm	Foreign Vlm	Retail NB	Foreign NB
	(1)	(2)	(3)	(4)	(5)
Hit 10%	16.0***	15.6***	0.22***	0.19***	−0.095***
	(0.21)	(0.21)	(0.01)	(0.03)	(0.02)
Hit 10% × Post	185.0***	183.9***	1.14***	1.87***	−0.90***
	(0.27)	(0.27)	(0.01)	(0.03)	(0.02)
Lagged Vlm	0.62***	0.61***	0.0057***		
	(0.00)	(0.00)	(0.00)		
Lagged Retail NB				0.33***	
				(0.00)	
Lagged Foreign NB					0.30***
					(0.00)
Hour FE	✓	✓	✓	–	–
Industry × Interval FE	–	–	–	✓	✓
Stock Controls	✓	✓	✓	✓	✓
Within R^2	.447	.447	.078	.104	.085
Observations	25,475,118	25,475,118	25,475,118	23,485,541	23,485,541

This table reports the coefficient estimates of the panel ordinary least squares (OLS) regression specification in equation (1). The different columns correspond to different dependent variables. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. Standard errors are clustered at the industry by interval level, and are reported in parentheses.

Challenges to Identification. The fundamental problem is that returns and trade flows are endogenously determined. Although the regression is specified so that the timing of returns, which determines the independent variable (VI occurrence or 10% breach), precedes that of the dependent variable (volumes and flows), there are still lingering concerns. I discuss three possibilities: news about fundamentals, liquidity shocks, and persistence in retail sentiment.

Suppose there is good news about the stock between the day's opening and the moment of static-up VI. Unless prices and flows adjust instantaneously, the prolonged effect of news will affect both returns from open and retail flow throughout the day. Liquidity shocks at large institutions work in the same way: if an institution is forced to make directional trades in a prolonged manner, this meta-order would affect both returns and retail flows. Persistence in retail sentiment is also similar: retail investors become excited about a stock, starts to put in aggressive buy orders, returns breach the threshold due to these orders, and excitement persists after the threshold breach.

The available remedies are (1) conditioning on prices rising by 10% relative to open and (2) controlling for preceding trade flows by different investors. Adding variables that precede the event, such as abnormal volume and liquidity measures, will further help control for news and liquidity shocks. However, we will see in the next subsection that the most concerning problem is retail sentiment *very close to* the VI thresholds.

Difference-in-Differences. In order to tackle the identification problem, I start by adapting the diff-in-diff setup proposed by Hautsch and Horvath (2019). The ideal experiment would be comparing two identical stocks that are under the influence of same news, liquidity shocks, and taste shocks, and only one stock randomly triggers a static VI. A reasonable alternative would be comparing two similar stocks that have both increased by 10%, while letting one stock randomly trigger a static VI. As does Hautsch and Horvath (2019), I face the same problem that the VI rules were introduced to all stocks with the same threshold. Thus, in this section I settle with the viable alternative of comparing a VI event with a pseudo-pause.

I first gather the static-up VI events for each stock i. Then, I keep the VI event if it is the first VI event of the day, then index it by $h = 1, \ldots, H$. There are 14,484 static-up VIs that satisfy this condition, so $H = 14,484$. For each VI event h, 20 minutes before and 60 minutes after the VI are sampled and 10-minute bins are constructed. These bins are assigned event time τ where $\tau = 0$ includes the 10 minutes immediately following the VI. Analogously, the first breaching of 10% is sampled from the pre-rule period and are given indices $h = H + 1, \ldots, H + K$. There are 9,089 such breaches so that $K = 9,089$. 20 minutes before and 60 minutes after these breaches are sampled in the same way.

Samples from actual VIs are considered treated so that $\text{Treated}_h = 1$ if $h \in \{1, \ldots, H\}$. Time is defined in terms of event time. The first difference is the difference between intervals with $\tau \geq 0$ and the interval $\tau = -2$. The second difference is between the VI and pseudo-pause events. Once we denote by D_τ the dummy variable for event time interval τ, we can write the diff-in-diff regression specification as follows:

$$y_{i\tau h} = \gamma \cdot \text{Treated}_h + \sum_{\substack{\tau=-2 \\ \tau\neq 0}}^{5} \gamma_\tau \cdot D_\tau + \sum_{\substack{\tau=-2 \\ \tau\neq 0}}^{5} \beta_\tau \cdot D_\tau \times \text{Treated}_h + \Gamma' \mathbf{X}_h + \alpha_i + \varepsilon_{i\tau h} \tag{2}$$

where $y_{i\tau h}$ is the outcome variable for stock i in event time τ for sample h, and α_i is the stock fixed effect. The outcome variable is either retail volume or net buy in each of the 10-minute interval τ. The controls are the same as in specification (1) and the coefficient of interest are β_τ's with $\tau \geq 0$. While γ_0 captures the effect of breaching the 10% threshold on the outcome variable,

β_0 captures the additional effect of triggering a static-up VI. The coefficients β_τ's are plotted in Figure 9.

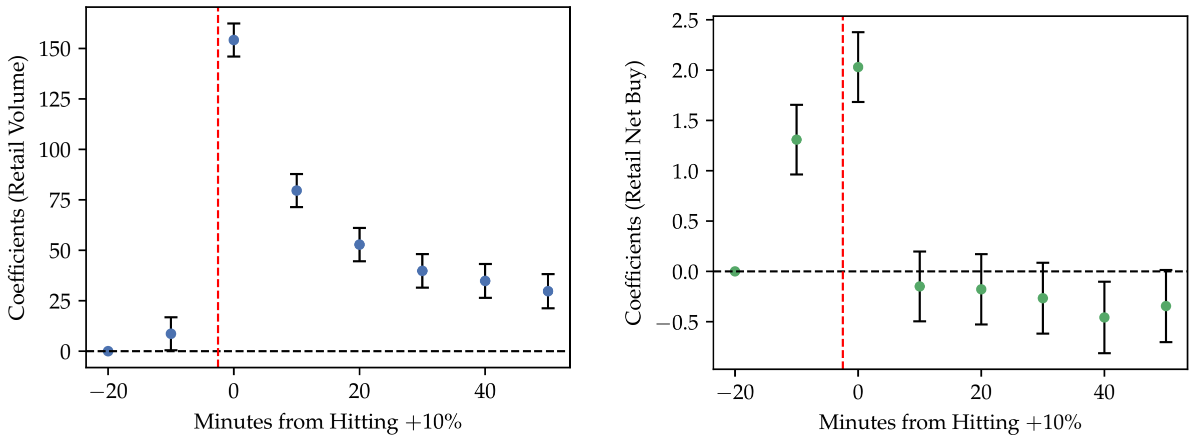

(a) Retail Volume

(b) Retail Net Buy

Figure 9. Retail Volume and Net Buy around 10% Breaches

This figure plots the estimates of the coefficients β_τ from the diff-in-diff regression specification in equation (2). 95% confidence intervals are displayed.

We see that the coefficient β_0 is positive and significant in both panels of Figure 9. This is consistent with previous subsections: retail investors increase trading abnormally, and they net buy stocks that just experienced a static-up VI, effectively betting on price continuation. The magnitudes of the coefficients are also largely unchanged relative to those of the panel OLS specification.

At the same time, we can observe a quite aggressive pre-trend in Figure 9b. This raises doubts about the necessary parallel trends assumption. One noticeable fact is that the pre-trend appears very abruptly around 10 minutes prior to the breaching of the threshold. This is a reason to believe that a sudden shift in retail sentiment is the more likely concern rather than news about fundamentals or institutional liquidity shocks. A fundamental news shock that is consistent with a static-up VI is *good news*. At the same time, it is reasonable to think that positive (negative) institutional (retail) net buy will be associated with this shock. This is not consistent with the pre-trend we observe. An explanation based on institutional liquidity shocks face this same issue. It also faces another issue that institutions are more likely to smooth their meta-orders throughout the day or week, while the pre-trend is consistent with a more sudden move.

Hautsch and Horvath (2019) attribute a similar pre-trend to the magnet effect. However, it is unnatural to consider the coefficient β_{-1} as the causal effect of an imminent trading pause that is yet to realize. The more natural interpretation would be: the fact that *a trading pause rule is in place* causes endogenous ex-ante reaction by traders during normal times and the *realization of trading pause* causes additional reaction by traders ex-post. This unfortunately means that $\tau < 0$

periods for pseudo-pause events and VI events are already dissimilar states.

3.3 Mechanisms

There are reasons to suspect that the heightened retail activity following VIs are due to the salience of the event among retail investors. As Figure 10 demonstrates, retail investors can get various alerts and filtered lists of stocks contingent on a stock's VI status through their mobile or desktop trading applications. It is also possible to filter on the proximity of the stock's price relative to the VI thresholds, which is an aspect more relevant to Section 4.

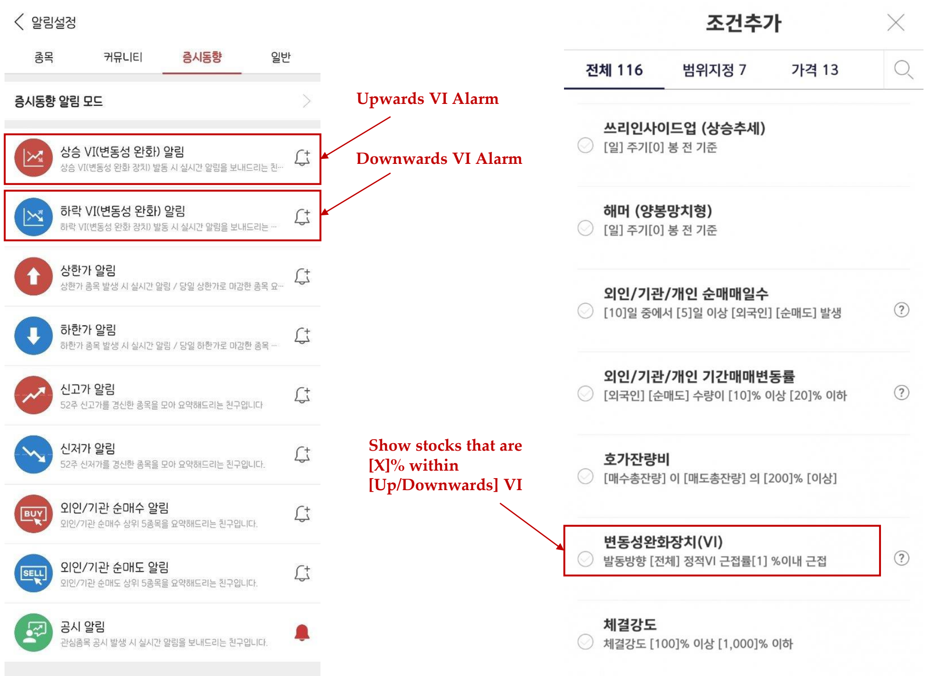

Figure 10. Features of Korean Mobile Trading Systems

These figures are screenshots taken from the mobile trading app of a popular Korean brokerage firm Kiwoom Securities. The figure on the left shows that a user can set an alarm when a VI is triggered for a chosen stock. The figure on the right shows that a user can put a filter to show only stocks whose prices are $X\%$ within a static-up or static-down VI.

Multiple VIs and Dissipation of Attention. In their study of trading halts in the Chinese market, Seasholes and Wu (2007) found a similar pattern of attention-induced trading for stocks that breach the upward price limits. Furthermore, they found that the effect becomes attenuated as the number of contemporaneous halts increases. If retail investors are limited in attention and trading capacity, this should also hold in my setting. I run another regression as specified in equation (2), but within different subsets of the pauses and pseudo-pauses. For instance, the blue dots in Figure 11 present the resulting coefficient estimates by using the VIs and pseudo-pauses that were the n-th event of the day, where $1 \leq n \leq 15$.

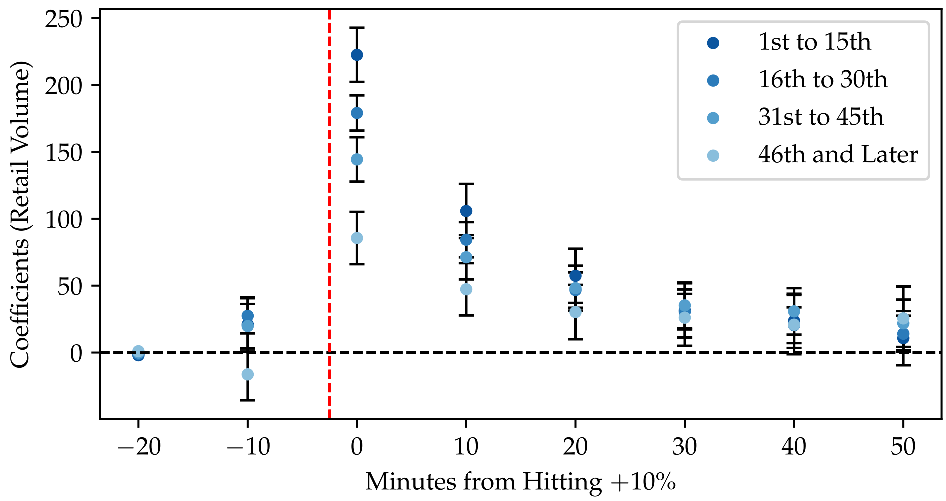

Figure 11. Retail Volume around 10% Breaches by Number of Previous Breaches

This figure repeats the plot in Figure 9a using four different subsets of the data. Both VIs and pseudo-pauses are ordered within the same day. The first group includes VIs and pseudo-pauses that are between the 1st and 15th to occur during their day of occurrence. Similarly, bins are assigned in increments of 15 occurrences. 95% confidence intervals are displayed.

In Figure 11, we see a clear monotone decrease in the coefficient estimates as we move onto subsets with late-coming VIs. This pattern supports the attention-based explanation. To address the concern that later VIs come at later times of the day, during which retail activity subsides, controls include time of the day dummies.

Preferred-Common Stock Pairs. To partially address the concern about reverse causality, I focus on a special case: preferred and common stock pairs. In Korea, around 100 established firms including Samsung, LG, and Hyundai, have floated preferred stocks on top of their common stocks. While the details differ slightly, preferred stocks offer a guarantee to pay extra dividends relative to their common stock counterparts at the cost of giving up voting rights. In theory, they should be subject to the same news about fundamentals and sentiment about the firm compared to their common stock counterparts. One notable aspect is that these preferred stocks have a much lower float, making them much more volatile relative to the common stocks. This means that there are many cases where a preferred stock breaches the 10% threshold while the common stock does not. I gather the *trades for the common stock* around *breaches of the preferred stocks*, for the pre- and post-rule periods, then run a diff-in-diff with these events. Because the common stock itself did not breach the 10% threshold, this exercise should be free from a taste shock causing the 10% breach and should only capture the attention spillover. Figure 12 shows the coefficient estimates, and we see effects for the retail volume, but lack power for retail net buy.

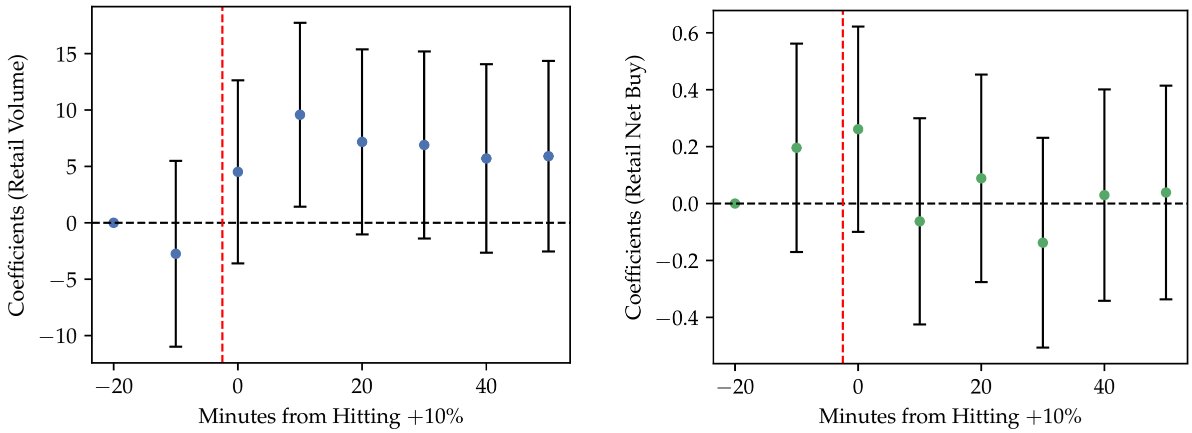

(a) Retail Volume

(b) Retail Net Buy

Figure 12. Retail Activity in Common Stocks around Preferred Stocks' VIs

This figure plots the estimates of the coefficients β_τ from the diff-in-diff regression specification in equation (2). 95% confidence intervals are displayed.

4 Ex-Ante Effect of VI Rules on Retail Trading

While the pre-trend in Section 3.2 poses difficulty in terms of identifying the causal effect of VI occurrence on retail trading, it is also what make the phenomenon more interesting. As the probability of triggering a static VI increases, traders appear to anticipate this possibility and alter their behavior (**?**). Thus, we should expect causal effect of rule adoption on ex-ante trading behavior as well.

4.1 Stylized Facts

I begin by presenting a drastic change in retail investor net buy behavior. First, I divide each stock i's time series on day d into 1-minute bins. Before moving on, I drop stock-minute observations that come after breaching of 10% for the pre-rule period and a static VI hit for the post-rule period. Essentially, all observations that are affected by the ex-post effects of threshold breaches are discarded. For each 1-minute interval $t \to t+1$, I use the price at the *beginning* of the interval to compute the return from that day's opening price:

$$r_{idt} = \frac{P_{idt}}{P^o_{id}} - 1$$

where P_{idt} is the price of stock i at the beginning of the interval $t \to t+1$ and P^o_{id} is the opening price of stock i on day d. Albeit imperfectly, this mitigates the concern that net trades in fact moved the prices to reach r_{idt}. I also compute the retail net buy during this same 1-minute interval, and call it nb_{idt}. If we see a high nb_{idt} associated with a high r_{idt}, this tells us that retail investors are net buying stocks that are expensive, or have appreciated a lot relative to the opening prices.

Mapping this to an illustration in Figure 13: if t is 9:05 a.m., r_{idt} will be 1.2% and nb_{it} will be the retail net buy during 9:05 and 9:06 a.m. Each stock-interval is assigned a return bin using the computed r_{idt}. The example interval in Figure 13 will be assigned to the [0%, 2%] bin if the bins were spaced out by 2%.

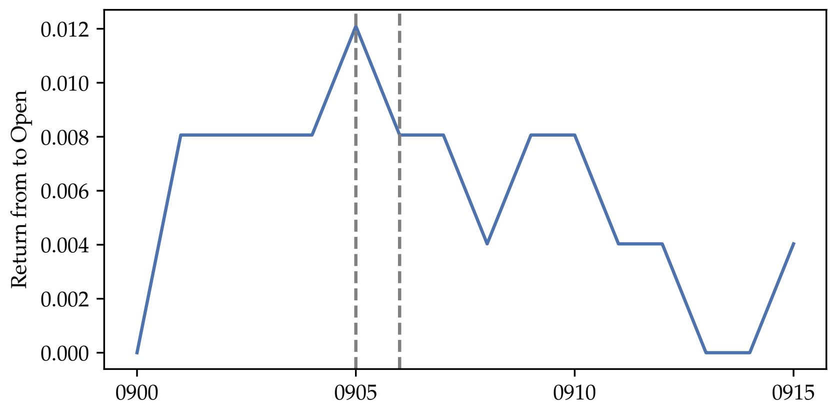

Figure 13. Return Bin Construction

This figure illustrates how 1-minute intervals are constructed and are assigned to return bins. The blue line represents a sample price path of a stock in terms of return from open. Return from open is defined as the return of the opening price of a 1-minute interval relative to the day's opening price.

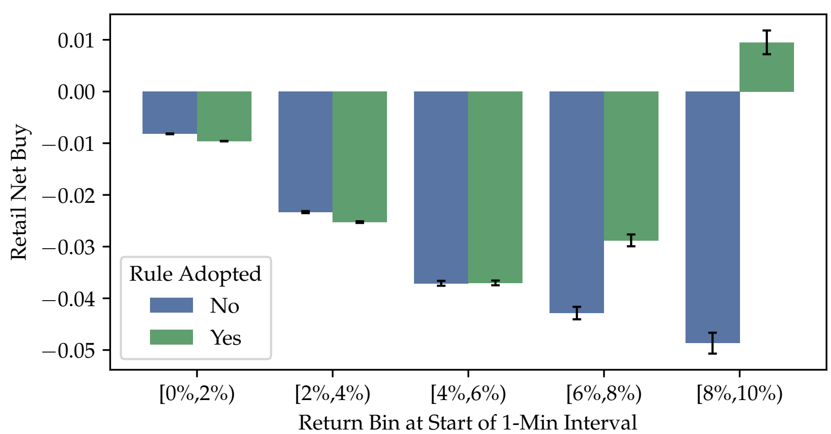

Figure 14. Average Retail Net Buy by Return Bins

This figure plots the average 1-minute net buy conditional on the return from open at the beginning of the 1-minute interval. The averages are computed for both pre-rule (blue bars) and post-rule (green bars) periods. 95% confidence intervals are constructed by bootstrapping.

For each return bin, the average retail net buy for the 1-minute intervals that fall in the return bin is computed. The results are shown in Figure 14. At medium to long horizons, retail investors provide liquidity to institutions. At the same time, institutions often make informed trades and do so by spreading out their orders. These in conjunction mean that on a day when institutions make informed buys, thereby pushing up prices, retail investors will end up net buying appreciated stocks. This narrative fits the pattern of the blue bars in Figure 14. In addition, disposition effect and mental accounting of retail investors are also consistent with the net selling of appreciated stocks.

What is surprising is the pattern of the green bars. After the VI rule adoption, they completely reverse their behavior near the static-up VIs and become net buyers of stocks that appreciated more than 8% intraday. Because all intervals that come after the breaching of $\pm 10\%$ were discarded, these net buy behaviors were affected by anticipation of possible VIs, but not by effects of realized VIs. Figure A11 plots the same quantities, but includes other bins.

I also look at the order book depth on the ask and bid sides conditional on the return bin of the current traded price. For these figures, I use a shorter sub-period, Jun 2014–May 2016, for which trade-level data is available. Here I find that Figure 15a shows the time-weighted average of order book depth on the ask side. Order book depth on either (bid or ask) side is computed by summing up the outstanding volume at the first three levels of the order book. It is then normalized by the average depth on both the ask and bid sides over the past 20 trading days. Thus, a value of 0.5 should be considered a *normal* relative depth. Time-weighted averages are computed conditional on the return bin of the latest traded price. This procedure is repeated separately for the pre-rule and post-rule periods.

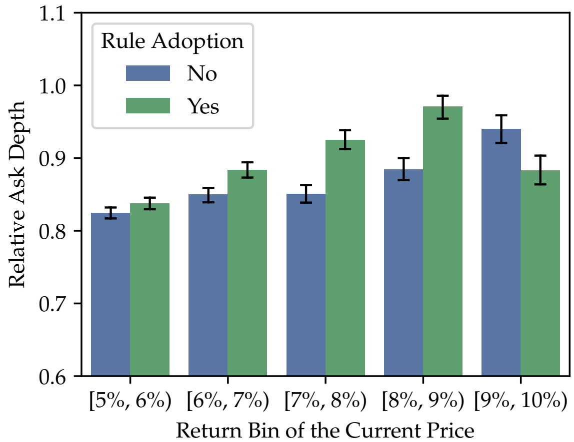

(a) Ask Depth

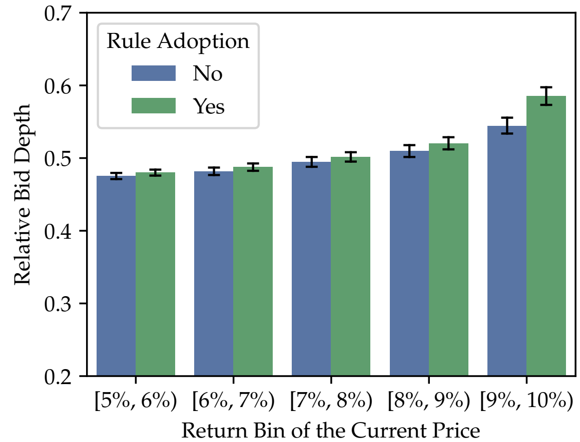

(b) Bid Depth

Figure 15. Relative Ask/Bid Order Depth by Return Bin

This figure plots the time-weighted average depth on the ask and bid sides conditional on the return bin of the current price. Ask (bid) depth is defined as the number of outstanding limit orders at the first three levels of the order book, normalized by the past 20 trading day average depth at both bid and ask combined. The heights of the blue (green) bars correspond to the pre-rule (post-rule) period averages. 95% confidence bands are constructed by bootstrapping.

In Figure 15a, we see that the ask depth conditional on the price being very close to the upper threshold is lower. One possibility is that traders postpone selling an appreciated asset when a VI is imminent. On the other hand, we see in Figure 15b that buying pressure near the threshold is higher in the post-rule period. These patterns are not consistent with attracting extra liquidity during extreme price upswings.

4.2 Empirical Results

I now test whether the patterns survive with various control variables. Also, the binning approach described in Figure 13 is an approximation. It works under the presumption that prices will lie in the same bin during the ensuing 1-minute interval. For instance, if prices start at 7.9% at the beginning of the minute and most of the trades in the minute take place around 9%, these trades will be mistakenly assigned to the [6%, 8%] return bin.

To get around this issue, I try a similar procedure with the trade-level data, albeit with a shorter sample period.[11] I follow the same procedure as before at the trade level. First, I discard the observations that come after a ±10% breach, then I assign return bins to each executed trade prices. The return bins are spaced by 1% instead for this analysis. For each stock-day, all retail buys and sells are summed up within the same bin, and the sums are normalized by shares outstanding. Consider a scenario where retail investors bought a total of 100 shares of stock

[11] Trade-level data spans June 2014–May 2016, while 1-minute frequency data continues to May 2018.

i during day d at an executed price (in terms of return from open) of 1.5% and sold a total of 100 shares at 2.5%. Suppose that there are 10,000 shares of stock i outstanding. Then, the retail buy volume for the bin [1%, 2%] is 100 bps, the retail sell volume of the bin [2%, 3%] is 100 bps, and the net buy for the two bins are $+100$ bps and -100 bps respectively. In such a way, stock-day-bin observations are constructed and used for the following regression:

$$y_{idb} = \gamma \cdot \text{Post}_{id} + \sum_{\substack{b=-9 \\ b \neq 0}}^{9} \gamma_b \cdot D_b + \sum_{\substack{b=-9 \\ b \neq 0}}^{9} \beta_b \cdot D_b \times \text{Post}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \varepsilon_{idb} \tag{3}$$

where the outcome variables are retail volume and retail net buy. Control variables are the same as the specification in equation (1).

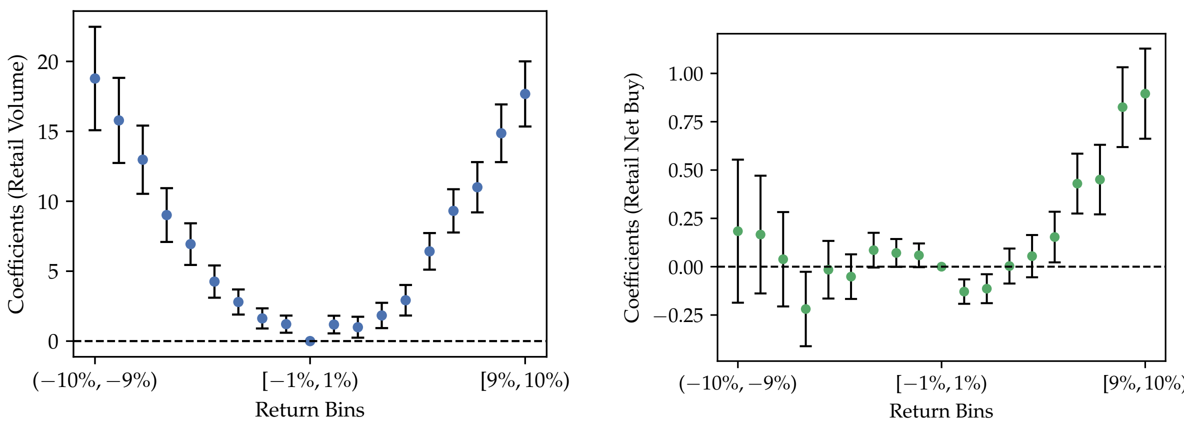

(a) Retail Volume

(b) Retail Net Buy

Figure 16. Difference in Ex-Ante Volume and Net Buy by Return Bins

This figure plots the estimates of the coefficients β_b from the diff-in-diff regression specification in equation (3). The outcome variable for panel (a) is retail volume and for panel (b) is retail net buy. 95% confidence intervals are displayed.

The resulting coefficient estimates for β_b's are shown in Figure 16. The scatter and confidence band at $x = b$ correspond to the bin [b%, $(b+1)$%]. Figure 16a shows that retail trading becomes more active around extreme prices. The regression controls for the effect of temporal changes through the Post_{id} term. We see that effects do not exist for moderate prices, but only for more extreme prices that are nearer the VI thresholds. Figure 16a captures the pattern documented in Figure 14: retail investors' tendency to net buy a stock increases when prices are near the static-up VI threshold. Patterns are similar if the trades are further normalized by duration of the prevailing trade prices, so that the outcome variables are in terms of intensity or rate of trades.

This pattern suggests a possible explanation of the pre-trend observed in Figure 9b. When we sample the static-up VIs ex-post, the price path is bound to pass through 8%, 9%, then 10%. The pattern just presented suggests that when prices are around these levels, retail investors are already net buying extensively and this will show up as a pre-trend in Figure 9b. This observation by no means extricates the reverse causality concern that aggressive retail buying near the threshold has a (reverse) causal effect on the realization of static-up VIs.

4.3 Trading Strategy

Cutoff Strategy. The presented explanation is that retail investors more aggressively buy a stock in anticipation of a static-up VI. Testing this statement directly requires an observable measure of retail investors' subjective probability of a static-up VI and their expectation of late-arriving buying pressure mentioned in Section 3. Because it is challenging to convincingly construct these measures, I resort to demonstrating an incentive to purchase that gets stronger as prices approach the upper threshold. I do this by presenting an extremely simple trading strategy: purchasing a stock as soon as its best ask price breaches X% and selling either (1) at the static-up VI auction, or (2) at the best bid once 15 minutes elapses, whichever happens first. This strategy may well be described as *bubble-riding* and similar ideas are also presented in Cho et al. (2003), Seasholes and Wu (2007), and Chen et al. (2019b).

The first thing to notice in Table 3 is that the strategy improves as the threshold X% increases from 8% to 9.5%. In other words, this is a strategy that improves as we buy at *more expensive* prices. The average return increases from 0.28% to 0.34% and standard deviation of returns decreases from 3.00% to 2.45%. Given that the sum of round trip brokerage fees and trading taxes amount to 0.3%, this is not profitable strategy in this crude form, especially considering the large standard deviation of returns. The trades are presented to demonstrate that due to the inflow of buy orders that come after the static-up VI, there is an incentive to buy at an already elevated price of 8% or 9%. Furthermore, this incentive becomes even stronger as we reach higher prices.

Table 3. Trading Strategy Returns

	Model	9.5%	9.0%	8.5%	8.0%
Count	1,960	6,475	7,210	8,053	9,071
Prob VI	0.70	0.84	0.75	0.64	0.55
Prob Profit	0.67	0.60	0.63	0.59	0.55
Mean	0.70	0.34	0.35	0.30	0.28
SD	3.29	2.45	2.67	2.85	3.00
Min	−12.23	−14.69	−14.59	−14.59	−18.48
25%	1.40	−0.36	−0.73	−1.35	−1.65
50%	1.41	0.38	0.68	0.82	0.64
75%	2.39	1.09	1.45	1.75	2.04
Max	22.40	22.40	22.40	22.40	22.40

This table reports the results from the trading strategy described in this subsection. The first column presents the result from using a classification model discussed in Appendix B. Other columns refer to results from a different buying cutoff. The top row specifies the number of cutoff breaches across all stocks during the period Jun 2015–May 2016. The second row shows the proportion of those breaches that end up in a static-up VI. The third row shows the proportion of trades that end up making a profit.

VI Prediction. The trade-off in choosing a higher cutoff is that the probability of exiting through a VI auction increases, but the purchase price also increases. If a trader can identify variables that predict a static-up VI, conditional on entering her trade at 8%, this naive strategy may be improved. I start by inspecting variables that may predict short-term retail activity to see if they can explain the occurrence of a VI, at least in-sample. If stock i breaches 8% on day d and time t, a trading position is opened. If this position exits through a VI auction, we let $y_{idt} = 1$, and let $y_{idt} = 0$ otherwise.

Daily covariates $\mathbf{X}_{id}$ include log market capitalization, Amihud measure, foreign investor holdings share, current day's overnight return (return from previous close to current day's open), past 20-day return, price deviation relative to 5-day moving average price, and log 5-day average trade volume. Intraday covariates $\mathbf{Z}_{idt}$ include log 5-minute trading volume, order imbalance (bid order volume over total order volume outstanding), and log order book depth. Using these variables, I run OLS regressions (linear probability model) and logistic regressions according to equations (4) and (5).

$$y_{idt} = \alpha + \Gamma'\mathbf{X}_{id} + \Lambda'\mathbf{Z}_{idt} + \varepsilon_{idt} \tag{4}$$

$$\mathbf{P}(y_{idt} = 1 \mid \mathbf{X}_{id}, \mathbf{Z}_{idt}) = \frac{\exp(\alpha + \Gamma'\mathbf{X}_{id} + \Lambda'\mathbf{Z}_{idt} + \varepsilon_{idt})}{1 + \exp(\alpha + \Gamma'\mathbf{X}_{id} + \Lambda'\mathbf{Z}_{idt} + \varepsilon_{idt})} \tag{5}$$

The results are shown in Table 4, where columns (1)–(3) show the coefficient estimates from

OLS and columns (4)–(6) show those from logistic regressions. Past returns, past volumes, and intraday variables are omitted from columns (1), (2), (4), and (5) to demonstrate the additional explanatory power of these variables.

The coefficient signs mostly line up well with intuition. We see that characteristics that indicate lower liquidity, such as smaller size, higher Amihud measure, and lower foreign institutional holding, lead to higher chance of static-up VI occurrences. Also, characteristics that indicate heightened sentiment, such as high overnight returns, higher recent trading volume, and higher order imbalance, also lead to higher chance of static-up VI occurrences. Unsurprisingly, the explanatory power of the regression models are still low because the exercise can be seen as attempting to predict future price movements.

While these coefficient estimates provide the intuition that imminent trading activity matters, these generalized linear models cannot capture high-dimensional relationships between VI occurrence and covariates. Because such relationships may provide extra predictive power, I train a tree-based model using the gradient boosting algorithm (Chen and Guestrin 2016) together with various other features. The details are reported in Appendix B. With this prediction model, I compute the trading results from entering positions at the 8% threshold only if an upward-VI is predicted. This result is reported in the first column of Table 3. We see that the average return is 0.70% and the probability of exiting through a VI is 0.70, while the standard deviation of returns is roughly similar as before (4th column). I offer these results to demonstrate that the foreseeable inflow of retail demand in case of a VI presents an incentive for going long—taking liquidity, rather than providing liquidity—near the upper thresholds.

Table 4. VI Prediction

	VI: Linear			VI: Logistic		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	2.58***	2.48***	2.56***	8.85***	8.37***	8.84***
	(0.15)	(0.16)	(0.16)	(0.67)	(0.71)	(0.73)
Log Market Cap	−0.08***	−0.09***	−0.09***	−0.35***	−0.38***	−0.38***
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Amihud	1.30***	1.40***	1.23***	6.00***	6.73***	6.00***
	(0.16)	(0.17)	(0.17)	(0.79)	(0.88)	(0.90)
Foreign Holdings	−0.34***	−0.31***	−0.30***	−1.61***	−1.46***	−1.45***
	(0.06)	(0.06)	(0.06)	(0.29)	(0.29)	(0.29)
Overnight Return		1.87***	1.89***		8.04***	8.26***
		(0.22)	(0.22)		(0.97)	(0.98)
Past 20-Day Return		0.01	0.01		0.05	0.03
		(0.02)	(0.02)		(0.11)	(0.11)
Deviation from 5-Day MA		0.11	0.07		0.46	0.29
		(0.07)	(0.07)		(0.31)	(0.32)
Log 5-Day Trading Volume		0.01**	0.01		0.05**	0.04
		(0.00)	(0.01)		(0.02)	(0.02)
Hours from Open			−0.02***			−0.08***
			(0.00)			(0.01)
Log 5-Minute Trading Volume			0.01***			0.04***
			(0.00)			(0.01)
Order Imbalance			0.10***			0.41***
			(0.02)			(0.10)
Log Book Depth			−0.02***			−0.07***
			(0.00)			(0.01)
Observations	11,375	11,375	11,375	11,375	11,375	11,375
R^2	0.05	0.06	0.07			

This table reports the estimates from OLS and probit regressions where the outcome variable is a dummy variable indicating whether the trade exited with a static-up VI. Standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. The sample includes all instances of stocks breaching 8% in terms of return from open during the period Jun 2016–May 2018.

5 Price Reversion and Retail Wealth Loss

Previous sections demonstrated retail investors' trading patterns around static VIs. A natural follow-up question is whether such patterns are profitable for them. This section demonstrates that it is not, at least for the retail investors as a whole. Data reveals that prices revert after both static-up and static-down VIs, and that this reversion is persistent in the medium term. This implies that retail investors are buying at the peak for static-up VIs and selling at the trough for static-down VIs. Foreign institutions take the other side of these retail trades. Retail investors effectively end up transferring wealth to foreign institutions through trading activity around VIs.

5.1 Price Reversion around Threshold Breaches

I start by plotting the prices, in terms of return from open, around $\pm 10\%$ breaches. For each stock i, prices at the end of each 1-minute interval t are recorded. These prices are converted into return from open as before. For the pre-rule period, the event time is defined in terms of minutes relative to the first t such that return from open breaches $\pm 10\%$ during interval $t \rightarrow t+1$. For the post-rule period, it is defined relative to the time of the first static VI of the day. 15 minutes preceding the breaches and 60 minutes following them are sampled. I drop pauses and pseudo-pauses that occur within the first 15 minutes of market open and last 60 minutes of market close, so that all event times can fall in the sampled observations. Because the resolution is at 1-minute frequencies, the closing price for the 0th minute is not always exactly equal to the 10% threshold. For instance, if prices fall from -9.5% to -11% within a 1-minute interval, the first sampled minute at event time 0 will have a price of -11%. This is the reason we see prices that are more extreme than $\pm 10\%$ for the blue lines in Figure 17.

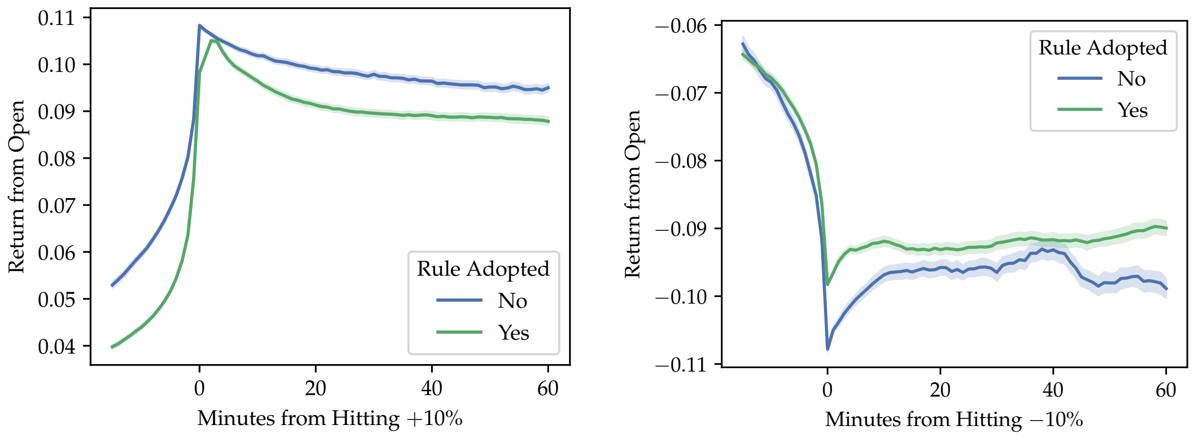

(a) Static-Up VI

(b) Static-Down VI

Figure 17. Prices around Static VIs

This figure plots average prices around VIs and pseudo-pauses, in terms of return from open. Return from open is computed using the closing price of each 1-minute interval according to Section 2.2. 95% confidence intervals are constructed by bootstrapping.

Figure 17 shows the results of the above procedure. Both Figures 17a and 17b show that prices revert after extreme returns from open. It should be noted that this pattern exists not just around the 10% thresholds, but at other less extreme thresholds such as 6% or 8%. The reversion is weaker at these cutoffs, however. If series of large flows impact prices and liquidity is scarce, we would expect to see such mean reversion if we sample ex-post at these threshold breaches. I do not attempt to explain the source of such reversion, but simply demonstrate that reversion around the VIs are extremely consistent as suggested by the tight confidence bands.

Recall that the first 2 minutes of the green lines are within the VI auction phase. This means that there are no executed prices during this brief period. This is why observations are missing for minute 1 in the figures above. Also, this means that the first observed prices at minute 2 are the clearing prices of the VI auctions. We can see that in both figures, prices tick up slightly for the green lines, indicating that the clearing prices of the VI auctions were slightly higher—about 0.5% on average—than the VI trigger prices.

5.2 Retail Wealth Losses

The price and retail trading patterns jointly suggest that retail investors are transferring wealth to foreign institutions around the VIs. I make back-of-the-envelope calculations for the magnitude of these losses by making the assumption that wealth is marked to market using the daily closing prices. A detailed explanation follows in the next subsections.

The bottom line is that during the two-year period following the rule adoption, retail investors transferred around an extra KRW 170 billion (approx. USD 150 million) to institutional investors on *extreme return stock-days*. A stock-day observation (i, d) during which the stock i breaches $\pm$10% is referred to as an extreme return stock-day. This loss happens over around 13,000 extreme return stock-days during this two-year period. This figure does not include the transfers from less sophisticated retail investors to other retail investors, which is expected to be at least an order of magnitude larger given the relative magnitudes of retail volume and retail net buy. Furthermore, total retail trading volume during these stock-days amount to KRW 210 trillion (approx. USD 190 billion). Due to round trip brokerage fees and trading taxes, which add up to at least 0.3% of trading volume, this costs retail investors an extra KRW 620 billion (approx. USD 550 million) over this period.

Losses around Static VIs. Suppose a static-up VI is triggered at minute t_0. Denote by x_t the number of shares net bought by retail during the interval $t \to t+1$ and y_t the net outflow of dollars due to these net trades. Because the exact holdings data is not available, we ignore the previous holdings and focus on the marginal change in marked-to-market wealth from t_0 onward. If P_{t_0+1} is the price at the end of this interval, retail investors' wealth change after the first interval will be

$$\Delta w_{t_0+1} = x_{t_0} P_{t_0+1} - y_{t_0}.$$

If prices are constant throughout the interval, $x_{t_0} P_{t_0+1} - y_{t_0}$ so that the wealth change is zero. Similarly, the marked-to-market wealth after T periods would be

$$P_{t_0+T+1} \sum_{t=0}^{T} x_{t_0+t} - \sum_{t=0}^{T} y_{t_0+t} \tag{6}$$

because over this period retail investors spent $\sum_{t=0}^{T} y_{t_0+t}$ dollars to accumulate $\sum_{t=0}^{T} x_{t_0+t}$ shares, each of which is valued at P_{t_0+T+1}. The implicit assumption is that the terminal price is the fair price, so later on sensitivity analysis needs to be done with respect to T.

All variables are directly observable in the data. Using the above definition, I first compute the loss in wealth for each static-up VI event using $T = 60$ minutes. I find that on average, retail investors lose approximately USD 9,000 per VI event and this amounts to approximately USD 45 million per year. Among these, 80% of the transfer is reaped by foreign institutions and the rest goes to domestic proprietary traders. Results for using different values for the fair price P_{t_0+T+1} are included in Table 5. The results indicate that estimated losses are larger if we use a larger T.

Table 5. Retail Trading Losses around Static-Up VIs

	60 Mins	120 Mins	180 Mins	Day' Close
Count	7,072	7,072	7,072	7,072
Mean	−10.09	−11.36	−12.41	−14.77
SD	80.55	86.57	100.52	114.95
Min	−2119.61	−2119.61	−3692.76	−3934.78
25%	−6.73	−7.44	−7.96	−9.33
50%	−0.48	−0.64	−0.69	−0.89
75%	1.94	1.79	1.83	1.73
Max	1146.77	991.34	991.34	1332.83

This table reports the retail trading losses computed according to equation (6) for different terminal prices. The last column uses the closing price of the day of the VI occurrence to mark-to-market. The reported values for mean, standard deviation, and percentiles are in terms of KRW millions.

Losses on Days with Extreme Returns. Using the procedure in equation (6), I can compute the profit/loss from trading in stock i on any given day d for each investor type. I assume for now that wealth is marked-to-market using stock i's closing price of day d, P_{id}. If retail investors in aggregate paid y_{id} in KRW terms throughout the day in order to accumulate x_{id} shares in net terms, then the profit/loss of the day is

$$\Delta w_{id} = P_{id} x_{id} - y_{id}. \tag{7}$$

I want to test if these daily net profits are lower on days that VIs occur. On a day in which a static-up VI occurs, retail investors will be net buying at a high price ($y_{id} \uparrow$) and will end up with more shares ($x_{id} \uparrow$) at the end of the day, but the price of these shares will have reverted ($P_{id} \downarrow$). I test this in a diff-in-diff set up where the first difference is between time relative to the VI rule adoption and the second difference is between normal and extreme return days. More specifically, I assign relative quarters, $q = -3, \ldots, 12$, to dates around the month of the rule adoption, and assign $\text{Treated}_{id} = 1$ if stock i breaches $\pm 10\%$ during day d. The first month after the rule adoption and the last month before the rule adoption are removed to allow for a burn-in period. Then, I run the following regression:

$$y_{idq} = \gamma \cdot \text{Treated}_{id} + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \gamma_q \cdot D_q + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \beta_q \cdot D_q \times \text{Treated}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \varepsilon_{idq} \tag{8}$$

where the outcome variables are daily retail volume and daily retail profit in KRW terms. The control variables are as in specification (3).

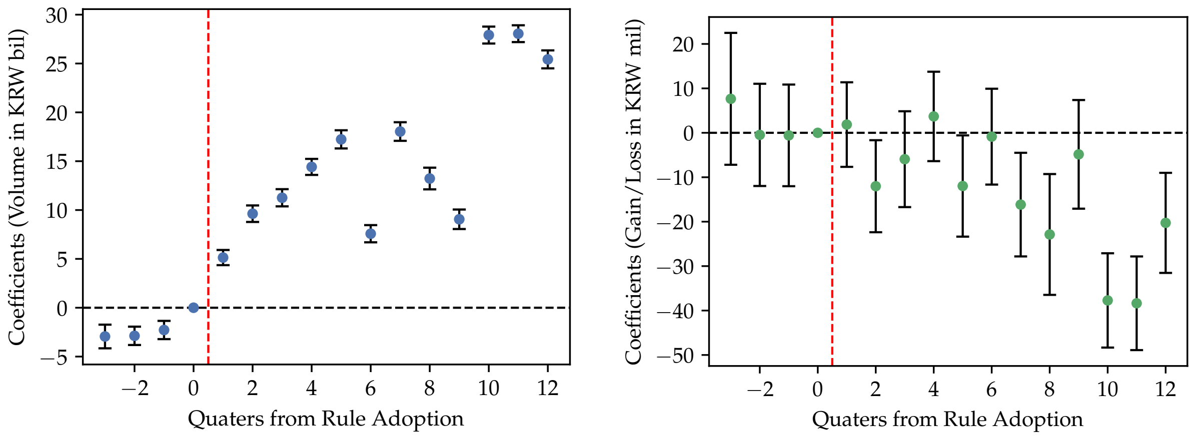

(a) Retail Volume

(b) Retail Profit/Loss

Figure 18. Retail Trading Volume and Profits on Normal vs. Extreme Return Days

This figure plots the estimates of the coefficients β_q from the diff-in-diff regression specification in equation (8). The outcome variables are retail trading volume in KRW billions and retail trading net profits in KRW millions. 95% confidence intervals are displayed.

The coefficients β_q's for $q > 0$ capture the effect of being in an extreme return day during post-rule quarters, after controlling for quarter fixed effects. Results in Figure 18b point towards larger losses, but are too noisy to interpret for most quarters. However, the results in Figure 18a show a clear pattern of increased activity on extreme return days. A coefficient estimate of 30 translates to a cost of: KRW 30 billion $\times$ 0.3% = KRW 90 million. This amounts to more than USD 80,000 of trading taxes and brokerage costs per stock-day. The costs add up to substantial amounts because there are more than 10,000 such stock-days in a year.

To gauge the average effect over the quarters adjacent to rule adoption, I run the following diff-in-diff regression:

$$y_{idq} = \gamma_1 \cdot \text{Treated}_{id} + \gamma_2 \cdot \text{Post}_{id} + \beta \cdot \text{Treated}_{id} \times \text{Post}_{id} + \Gamma' \mathbf{X}_{id} + \alpha_i + \gamma_q + \varepsilon_{idq} \tag{9}$$

where $\text{Treated}_{id} = 1$ if stock i breached $\pm 10\%$ during day d and $\text{Post}_{id} = 1$ if day d falls in the post-rule period. The stock-day-level control variables are lagged by 20 trading days, except for the past return variables which are lagged by 1 trading day.

Table 6. Retail Trading Volume and Losses on Extreme Return Days

	Retail Volume		Retail Profit	
	(1)	(2)	(3)	(4)
Constant	-77.89^{***}	-133.87^{***}	73.53^{***}	114.56^{***}
	(0.34)	(20.78)	(3.75)	(30.37)
Hit ±10%	10.92^{***}	10.59^{***}	-20.33^{***}	-19.51^{***}
	(0.21)	(0.95)	(2.26)	(1.71)
Hit ±10% × Post	16.88^{***}	16.75^{***}	-12.08^{***}	-13.10^{***}
	(0.23)	(2.65)	(2.53)	(4.40)
Log Market Cap	2.91^{***}	5.10^{***}	-2.77^{***}	-4.37^{***}
	(0.01)	(0.78)	(0.13)	(1.14)
Amihud	4.36^{***}	10.84^{***}	-8.79^{**}	-10.36^{***}
	(0.36)	(2.86)	(3.93)	(3.22)
Past Volatility	125.41^{***}	79.88^{***}	-39.10^{**}	-10.75
	(1.60)	(10.52)	(17.63)	(15.33)
Market Beta	0.63^{***}	0.06	-1.11^{***}	-0.74
	(0.03)	(0.31)	(0.32)	(0.54)
Past 20-Day Return	19.51^{***}	19.49^{***}	-4.86^{***}	-4.55^{**}
	(0.11)	(1.92)	(1.26)	(2.27)
Past 200-Day Return	2.22^{***}	1.98^{***}	-1.61^{***}	-1.23^{**}
	(0.03)	(0.69)	(0.33)	(0.57)
Stock FE	–	✓	–	✓
Quarter FE	–	✓	–	✓
Observations	797,293	797,293	797,293	797,293
R^2	0.20	0.18	0.002	0.001

This table reports the estimates from diff-in-diff regressions as specified in equation (9) where the outcome variables are daily retail trading volume and daily retail trading profit. Trading volume is in KRW billions and trading profit is in KRW millions. Standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% level, respectively. For columns (2) and (4), standard errors are two-way clustered at the stock and quarter levels.

The coefficient of interest is β, which is reported in the third row of Table 6. We see from the second row that retail investors trade more extensively and also make larger losses on extreme return days. The third row shows that both of these effects intensify during the post-rule period. These results are inconsistent with the intention of cooling-off markets in case of extreme price movements.

One concern may be that extreme returns, which determines Treated_{id}, is in fact a result of high trading volume y_{idq}. This can be addressed to some degree by taking the trading activity that comes *after* the occurrence of ±10% breaches. Because this necessarily drops all preceding trades, average trading activity has to be used instead of daily gross trading activity in order to make a fair comparison. This procedure does not affect the significance of coefficient estimates, but complicates the interpretation.

5.3 Effect on Overall Market Conditions

Despite these concerning patterns, circuit breaker-like rules are ubiquitous. Sifat and Mohamad (2019) count 48 stock exchanges with trading halts, 98 with price limits, and 31 with volatility interruption mechanisms out of the 152 exchanges studied, as of 2018. More than 85% of world's stock market capitalization is subject to some form of stock-level trading interruption rules. This leads us to expect that such rules improve price stability and liquidity. I start by testing whether this is true. The difficulty arises from the fact that the rule started to apply to all stocks with the same thresholds on the same day. This precludes the usual empirical designs related to natural experiments since we lack a control group.

Volatility as Treatment Dosage. While there were no variation in the threshold level, I propose a way to turn this nuisance on its head. The threshold for a static VI, namely 10%, was chosen arbitrarily. For a highly volatile stock with daily volatility of 5%, this threshold is *tight*, whereas for a stock with daily volatility of 1%, it is *loose*. In this sense, stocks were treated to a different degree: volatile stocks received a higher dosage of treatment. For instance, Samsung Electronics' common stock never triggers a static VI during the sample period used here.

In the benchmark specification, stocks are sorted by their daily return volatility during the past 60 trading days. Then, the top quintile is considered the treatment group while the bottom quintile is considered the control group. The other three quintiles are discarded. The results remain similar in different variations—e.g., including the other quintiles or using terciles instead. Also, results are similar when sorts are done based on past 20 quarters before the beginning of the pre-rule period.

Effect on Volatility. Outcome variables related to a stock's volatility are: 5-minute realized volatility, daily return volatility, high-low price range, and number of $\pm 10\%$ breaches. Daily return volatility is computed at the monthly frequency using daily close-to-close returns. 5-minute realized volatility for a given day is computed by summing squared 5-minute log returns within each day, then taking the square root of this number. Then, this daily value is averaged at the monthly frequency. High-low price range is computed by taking the difference between the daily high and daily low, and then dividing this number by the opening price. Again, this is averaged at the monthly frequency. The number of $\pm 10\%$ breaches is computed by counting the number of days during which a stock breaches $\pm 10\%$ relative to its opening price.

Event study plots in Figure 19 show the coefficients β_q's from running the regressions under the following specification:

$$y_{ijmq} = \gamma \cdot \text{Treated}_{im} + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \gamma_q \cdot D_q + \sum_{\substack{q=-3 \\ q \neq 0}}^{12} \beta_q \cdot \text{Treated}_{im} \times D_q + \Gamma' \mathbf{X}_{ijmq} + \alpha_j + \varepsilon_{ijmq} \tag{10}$$

where y_{ijmq} is the outcome variable for stock i in month m, α_j is the industry fixed effect, $\text{Treated}_{im} = 1$ if stock i falls in the fifth quintile in terms of past volatility, and D_q is a dummy variable that is equal to 1 if the observation falls in quarter q. Logarithms are taken for the outcome variables to test for relative changes. Quarters are relative to the rule adoption and the month of the rule adoption, June 2015, is dropped. Vector of controls $\mathbf{X}_{ijmq}$ include log market capitalization, market beta using past 120 trading days, and past 5-day, 20-day, 200-day returns. More importantly it also includes quintile-specific slopes for market volatility z_m and market return x_m in month m:

$$\delta_1 \cdot z_m + \delta_2 \cdot \text{Treated}_{im} \times z_m + \delta_3 \cdot x_m + \delta_4 \cdot \text{Treated}_{im} \times x_m.$$

This is to partially address the inherent limitation of this specification, which is that selection on past volatility happens by design. Since more volatile stocks would respond differently to market conditions, heterogeneous slopes are added. Suppose, for example, we take the simplest factor model

$$r_{it} = \alpha_i + \beta_i r_{mt} + \varepsilon_{it}$$

where r_i is stock i's return and r_m is the market return. Then, as long as the idiosyncratic components are small, the difference between volatility of treated and untreated stocks amplified during volatile times. The specification in equation (10) includes the market return and market volatility as controls to partially address this issue, but nonlinear dependencies cannot be addressed with the current specification.

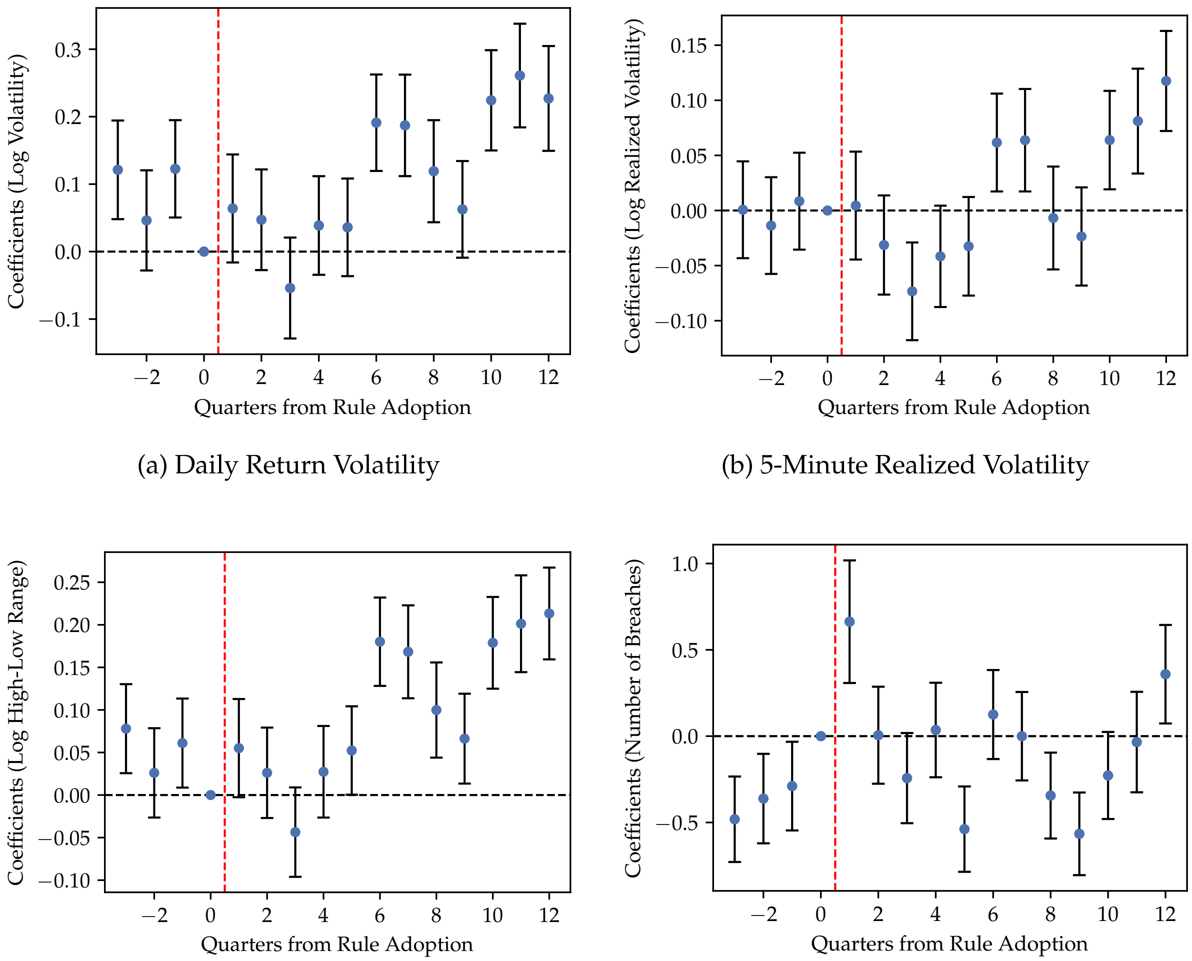

(a) Daily Return Volatility

(b) 5-Minute Realized Volatility

(c) High-Low Price Range

(d) Number of ±10% Breaches

Figure 19. Measures of Volatility and Rule Adoption

This figure plots the estimates of the coefficients β_q's from the diff-in-diff regression specification in equation (10). Logs are taken for the outcome variables in panels (a), (b), and (c). 95% confidence intervals are displayed.

Coefficients displayed in panels (a), (b), and (c) represent the average *differences in proportional changes* because outcome variables are log-transformed. The expected policy effect would be either decreased volatility or decreased frequency of extreme returns—i.e., negative and significant coefficients. The event study plots are not consistent with this policy effect.

I also run the following regression specification that is closer to the usual diff-in-diff specification:

$$y_{ijmq} = \gamma \cdot \text{Treated}_{im} + \beta \cdot \text{Treated}_{im} \times \text{Post}_m + \Gamma' \mathbf{X}_{ijmq} + \alpha_j + \gamma_q + \varepsilon_{ijmq} \tag{11}$$

where $\text{Post}_m = 1$ if month m comes after the rule adoption date, α_j and γ_q are stock and quarter fixed effects respectively, and all other variables are the same as before. Both quarter and industry fixed effects are included. Here I use the 4 quarters around the rule adoption instead.

Table 7. Diff-in-Diff Coefficient Estimates for Volatility

	Return Vol	Realized Vol	High-Low Range	Num Breaches
	(1)	(2)	(3)	(4)
Treated × Post	-0.02	-0.03^{*}	-0.01	0.30
	(0.03)	(0.02)	(0.02)	(0.19)
Stock Controls	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	7,558	7,558	7,558	7,558
Within R^2	0.38	0.52	0.54	0.22

This table reports the coefficient estimates of the diff-in-diff regression specification in equation (11). Logs are taken for the outcome variables in columns (1), (2), and (3). Standard errors are two-way clustered at the industry and quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Table 7 reports the coefficients β's from running the regressions according to equation (11). We see that only the coefficient in column (2) is weakly negative, and that the point estimate in column (4) is in fact positive. The introduction of static VI does not appear to depress the price variability in the treatment group.

Effect on Liquidity. I repeat the exercise in the previous subsection with measures of liquidity. Outcome variables for liquidity are: relative bid-ask spread, order book depth, and the Amihud measure. The relative bid-ask spread is computed by normalizing the bid-ask spread by the mid-price, and then averaging this number at the monthly frequency. Order book depth is the monthly average of outstanding total volume at the first three levels of the order book on both the ask and bid sides. The Amihud measure is computed at the monthly frequency as in Amihud (2002), except volume is expressed in KRW billion.

Table 8. Diff-in-Diff Coefficient Estimates for Liquidity

	Bid-Ask Spread	Book Depth	Amihud Measure
	(1)	(2)	(3)
Treated × Post	−0.04	0.05	−0.14*
	(0.03)	(0.10)	(0.07)
Stock Controls	✓	✓	✓
Industry FE	✓	✓	✓
Quarter FE	✓	✓	✓
Observations	7,239	7,239	7,239
Within R^2	0.31	0.08	0.75

This table reports the coefficient estimates of the diff-in-diff regression specification in equation (11). Logs are taken for the outcome variables. Standard errors are two-way clustered at the industry and quarter level, and are shown in parentheses. *, **, and *** denote 10, 5, and 1% statistical significance respectively.

Again, coefficients displayed in Table 8 represent differences in proportional changes because outcome variables are log-transformed. If the policy works as expected, we should see negative coefficients in columns (1) and (3), and a positive coefficient in column (2). Only the coefficient in column (3) is consistent with the policy objectives. However, the coefficient is statistically significant only at the 10% level.

Preferred Stocks as Treatment Group Using the same reasoning as before, I use preferred-common stock pairs as treatment-control pairs: preferred stocks are more volatile and can be considered to face a tighter threshold. In this case, adding firm fixed effects mitigates the selection problem in the previous section. I repeat the regression in equation (10) with underlying firm fixed effects instead of industry fixed effects.

Table 9. Diff-in-Diff Coefficient Estimates for Volatility

	Return Vol	Realized Vol	High-Low Range	Num Breaches
	(1)	(2)	(3)	(4)
Treated × Post	-0.02	0.05	0.07	0.48
	(0.10)	(0.06)	(0.08)	(0.46)
Firm FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	3,526	3,526	3,526	3,526
Within R^2	0.02	0.07	0.04	0.02

This table reports the coefficient estimates of the diff-in-diff regression specification in equation (11). Standard errors are clustered at the quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

Table 10. Diff-in-Diff Coefficient Estimates for Volatility

	Bid-Ask Spread	Book Depth	Amihud Measure
	(1)	(2)	(3)
Treated × Post	-0.19^{***}	-0.25^{***}	-0.68^{***}
	(0.05)	(0.08)	(0.22)
Firm FE	✓	✓	✓
Quarter FE	✓	✓	✓
Observations	3,379	3,379	3,379
Within R^2	0.58	0.69	0.76

This table reports the coefficient estimates of the diff-in-diff regression specification in equation (11). Standard errors are clustered at the quarter level, and are shown in parentheses. *, **, and *** denote 10%, 5%, and 1% statistical significance respectively.

The results are presented in Table 9 and Table 10. As before, we do not see improved price stability in the treatment group. In fact, we see positive, but insignificant point estimates in columns (2), (3), and (4) of Table 9. In Table 10, we see improvements in terms of bid-ask spread and the Amihud measure. However, the coefficient estimate in column (2) is negative, which is contrary to the expected effect of the rule introduction. Again, results are at best mixed, and it is difficult to make definitive statements on the efficacy of VI rule adoption on overall market conditions.

6 Conclusion

Using data from the Korean stock market, I present a series of unintended consequences of stock-level trading pauses, whose stated goal is "investor protection." During the years following rule introduction, retail investors as a whole have transferred hundreds of millions of dollars of wealth to foreign institutions on extreme return days which triggered trading pauses. These losses emerge due to the tendency of retail investors to bet on price continuation around these pauses. They have lost even more wealth through trading taxes and costs related to excessive attention-induced trading caused by the salience of this mechanism. Furthermore, this paper finds scant evidence of improvements in market conditions, such as volatility or liquidity, that benefit these retail investors.

The Korean stock market proves to be particularly well-suited in documenting these empirical facts due to its institutional features related to trade flow data. First, investor type-level flow data exists for a relatively long time series at a high resolution, and spans both the pre- and post-rule periods. Second, these equity flows are exhaustive because KRX is the sole exchange through which all stock transactions happen. Lastly, retail investors actively, and directly, participate in the stock market, so that clear patterns emerge around extreme price movements even at high frequencies. These edges allow me to quantify the retail wealth losses that occur immediately around trading pause events, and on days during which pauses are triggered.

Although such uniqueness is what enables the project, the lessons we can draw are broader. Previous studies on trading pauses and halts, whose settings span all major markets in the US, China, and Europe, hint at the universality of post-pause market instability. The similarities lead us to suspect that retail investors are being left vulnerable in other major markets as well. A more general takeaway is that preponderance of behavioral actors may derail even simple regulatory interventions.

There are a couple of venues that this paper can be pushed further. By relying on the arbitrariness of the pause thresholds, I attempt to circumvent the identification challenge posed by the fact that the rule was adopted to all stocks on the same day. However, this remains an imperfect alternative. Working directly with the KRX to carry out relevant randomized control experiments would address identification concerns. Also, explicitly modeling and estimating the proportion of inattentive arbitrageurs and inattentive behavioral actors will allow us to think about counterfactual policies that can better achieve the original regulatory goals.

References

Amihud, Y. (2002). Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. *Journal of Financial Markets*, 5(1):31–56.

Bali, T. G., Cakici, N., and Whitelaw, R. F. (2011). Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns. *Journal of Financial Economics*, 99(2):427–446.

Barber, B. M., Huang, X., Odean, T., and Schwarz, C. (2022). Attention-Induced Trading and Returns: Evidence from Robinhood Users. *The Journal of Finance*, 77(6):3141–3190.

Barber, B. M., Lee, Y.-T., Liu, Y.-J., and Odean, T. (2009). Just How Much Do Individual Investors Lose by Trading? *The Review of Financial Studies*, 22(2):609–632.

Barber, B. M. and Odean, T. (2008). All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *The Review of Financial Studies*, 21(2):785–818.

Barrot, J.-N., Kaniel, R., and Sraer, D. (2016). Are Retail Traders Compensated for Providing Liquidity? *Journal of Financial Economics*, 120(1):146–168.

Boehmer, E., Jones, C. M., Zhang, X., and Zhang, X. (2021). Tracking Retail Investor Activity. *The Journal of Finance*, 76(5):2249–2305.

Chen, H., Petukhov, A., and Wang, J. (2019a). The Dark Side of Circuit Breakers. page 57.

Chen, T., Gao, Z., He, J., Jiang, W., and Xiong, W. (2019b). Daily Price Limits and Destructive Market Behavior. *Journal of Econometrics*, 208(1):249–264.

Chen, T. and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA. Association for Computing Machinery.

Cho, D. D., Russell, J., Tiao, G. C., and Tsay, R. (2003). The Magnet Effect of Price Limits: Evidence from High-Frequency Data on Taiwan Stock Exchange. *Journal of Empirical Finance*, 10(1):133–168.

Christie, W. G., Corwin, S. A., and Harris, J. H. (2002). Nasdaq Trading Halts: The Impact of Market Mechanisms on Prices, Trading Activity, and Execution Costs. *The Journal of Finance*, 57(3):1443–1478.

Engelberg, J., Sasseville, C., and Williams, J. (2012). Market Madness? The Case of "Mad Money". *Management Science*, 58(2):351–364.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer, New York, NY.

Hautsch, N. and Horvath, A. (2019). How Effective Are Trading Pauses? *Journal of Financial Economics*, 131(2):378–403.

Hirshleifer, D. A., Myers, J. N., Myers, L. A., and Teoh, S. H. (2008). Do Individual Investors Cause Post-Earnings Announcement Drift? Direct Evidence from Personal Trades. *The Accounting Review*, 83(6):1521–1550.

Kim, J. and Kim, M. (2022). Behavioral Biases and the Trading of Individual Investors in the Korean Stock Markets. Technical report, Korea Capital Market Institute.

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

Lucas, R. E. (1976). Econometric Policy Evaluation: A Critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46.

Luo, C., Ravina, E., Sammon, M., and Viceira, L. M. (2022). Retail Investors' Contrarian Behavior Around News, Attention, and the Momentum Effect.

Odean, T. (1999). Do Investors Trade Too Much? *American Economic Review*, 89(5):1279–1298.

Scheinkman, J. A. and Xiong, W. (2003). Overconfidence and Speculative Bubbles. *Journal of Political Economy*, 111(6):1183–1220.

Seasholes, M. S. and Wu, G. (2007). Predictable Behavior, Profits, and Attention. *Journal of Empirical Finance*, 14(5):590–610.

Shefrin, H. and Statman, M. (1985). The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *The Journal of Finance*, 40(3):777–790.

Sifat, I. M. and Mohamad, A. (2019). Circuit Breakers as Market Stability Levers: A Survey of Research, Praxis, and Challenges. *International Journal of Finance & Economics*, 24(3):1130–1169.

Subrahmanyam, A. (1994). Circuit Breakers and Market Volatility: A Theoretical Perspective. *The Journal of Finance*, 49(1):237–254.

Yan Du, D., Liu, Q., and Rhee, S. G. (2009). An Analysis of the Magnet Effect Under Price Limits. *International Review of Finance*, 9(1-2):83–110.

Appendix A Data Description

Table A1. Available Variables from KRX

Data Category	Variables
1-Minute Frequency Price Data	Open, High, Low, Close, Trading volume, Trading value, Accumulated trading volume, Accumulated trading value, Best bid order price, Best ask order price, Residual quantity at best bid order, Residual quantity at best ask order, Bid-ask midpoint, Depth-weighted average limit price, Total cancelled IOC, Total cancelled FOK, Cancelled order ratio, Cancelled order volume ratio, Buy-sell imbalance in number of trades, Buy-sell imbalance in trading volume, Order imbalance in number of orders, Order imbalance in order volume, Average quoted spread, Average percentage quoted spread, Average effective spread, Average realized spread, HS adverse selection costs, Depth in volume, Depth in number of orders
1-Minute Frequency Trade Flow Data	Investor code, Accumulated buy trading volume, Accumulated buy trading value, Accumulated sell trading volume, Accumulated sell trading value
Trade-Level Data	Expected matching price, Expected matching quantity, Total residual quantity at ask order, Ask step 1-10 price, Ask step 1-10 residual quantity, Total residual quantity at bid order, Bid step 1-10 price, Bid step 1-10 residual quantity, Order trade acceptance no., Accumulated trading volume by market, Accumulated trading value by market, Aggregated residual quantity at step1-10 best ask order, Aggregated residual quantity at step1-10 best bid order, Last Step 1 best ask order price, Last Step 1 best bid order price, Total residual quantity at last ask orders, Total residual quantity at last bid orders, etc.
Other Market Data	Short volume, Short interest, Book-to-market ratio, Dividend yield, Shares outstanding, Foreign holdings

This table lists the available variables for each dataset. More details are available on the KRX website.

Appendix B Trading Strategy with a Tree-Based Classifier

In this section, I provide the details of the training procedure of the gradient boosting algorithm that predicts imminent VIs. Gradient boosting is a classification (and possibly a regression) technique that employs weak decision tree classifiers. The individual tree predictions are aggregated to form an ensemble classifier. I use XGBoost (Chen and Guestrin 2016) to implement the algorithm.

The sample consists of stock-time observations when the stock first breaches 8% return relative to its opening price. This is a classification problem whose outcome variable $y = 1$ if this stock triggers a static-up VI within the next 15 minutes and $y = 0$ otherwise.

There are two limitations to the data I work with. The first is that trade-level data is only available for the period June 2014–May 2016. In order to compute realistic trading performance, I assume that purchases happen at the ask and sales happen at the bid. If I use only the trade-level data to train and test the model, I am forced to work with less than 10,000 data points. I choose the alternative of training the data using 1-minute frequency data, which span June 2014–May 2018, and using the trade-level data only when computing the trading performance. I perform cross-validation on the data from June 2016–May 2018, without using any of the trading period, and also train the final model on this same period afterwards. This unfortunately means that the classification model uses data from the future. Still, none of the trading period data is used in the training and cross-validation process.

The second limitation is the relative dearth of data. The above sampling procedure yields around 15,000 to be used in training and cross-validation, and around 9,000 trading instances. Because a tree-based method is chosen to capture high-dimensional relationship between the features, this is a small sample. One limiting factor comes from the pool of chosen stocks. To be consistent with the main analyses of the paper, much of the small-cap stocks are discarded. Because these are the stocks that most often satisfy the sampling criterion, I end up with a much smaller sample.

First, features are constructed according to the definitions in Table A2. Then, 5-fold cross-validation is used to compute the average out of sample prediction performance for each hyperparameter combination. As with other machine learning algorithm, the gradient boosting algorithm is prone to overfitting and most of the hyperparameters control the amount of regularization imposed. Baseline values for tree depth, number of estimators, and learning rate were taken from Hastie et al. (2009).

Table A2. Features and Definitions

Variable Name	Definition
Log Market Capitalization	log(shares outstanding $\times$ day's closing price), lagged by 5 days
Market Beta	Beta w.r.t. KOSPI, lagged by 5 days
Amihud Measure	$10^9 \times \frac{\lvert\text{day's return}\rvert}{\text{day's volume in KRW}}$, lagged by 5 days
Foreign Holdings	Percentage of foreign investor holdings, lagged by 5 days
Overnight Return	Return from previous day's close to the trading day's open
5-Day Return	Closing price return in the past 5 days, lagged by 1 day
20-Day Return	Closing price return in the past 20 days, lagged by 1 day
200-Day Return	Closing price return in the past 200 days, lagged by 1 day
5-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 5 days}} - 1$
20-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 20 days}} - 1$
40-Day Disparity	$\frac{\text{buy price}}{\text{average closing price in the past 40 days}} - 1$
Disparity from 5-Day High	$\frac{\text{buy price}}{\text{past 5-day high price}} - 1$
Disparity from 20-Day High	$\frac{\text{buy price}}{\text{past 20-day high price}} - 1$
Disparity from 200-Day High	$\frac{\text{buy price}}{\text{past 200-day high price}} - 1$
5-Day Trading Volume	Trading volume in KRW in the past 5 days
5-Day Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 5 days}}{\text{trading volume in KRW in the past 40 days}}$
10-Day Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 10 days}}{\text{trading volume in KRW in the past 40 days}}$
Opening Trading Volume	Trading volume during the first minute of the trading day
Normalized Opening Trading Volume	$\frac{\text{opening trading volume}}{\text{40-day average of opening trading volume}}$
5-Day Normalized Opening Volume	$\frac{\text{5-day average of opening trading volume}}{\text{40-day average of opening trading volume}}$
10-Day Normalized Opening Volume	$\frac{\text{10-day average of opening trading volume}}{\text{40-day average of opening trading volume}}$
5-Minute Normalized Trading Volume	$\frac{\text{trading volume in KRW in the past 5 minutes}}{\text{trading volume in KRW in the past 40 days}}$
Order Imbalance	$\frac{\text{bid volume outstanding in order book}}{\text{bid volume outstanding in order book} + \text{ask volume outstanding in order book}}$
Order Book Depth	Sum of volume at the first bid and first ask
Hours from Open	Hours relative to market open

This table lists the variables used for VI prediction. *Trading day* refers to the day that the stock in question breaches 8%. *Buy price* refers to the trading day's opening price $\times$ 1.08.

Preliminary analyses on performance were run as shown in Figure A1 to come up with the hyperparameter grid/region. Then, a grid search is performed over the combinations of hyperparameters listed on Table A3.

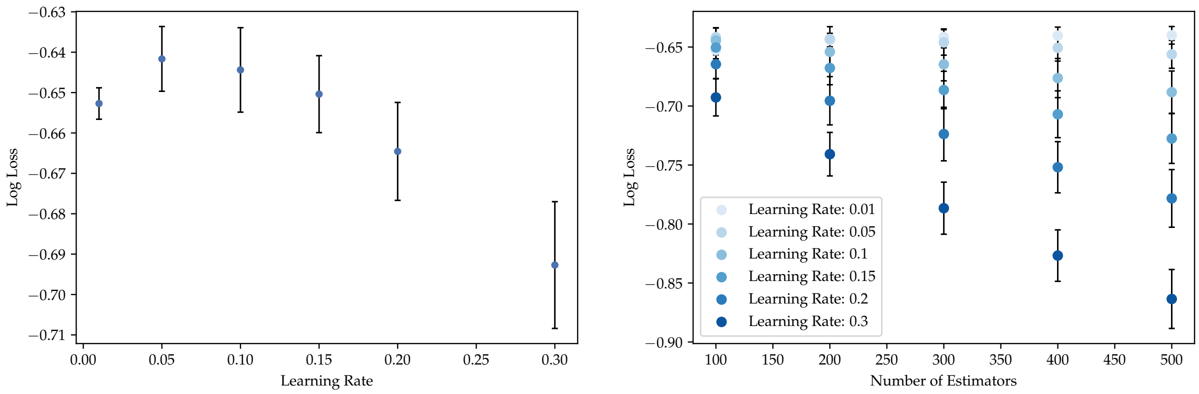

(a) Learning Rate

(b) Number of Estimators and Learning Rate

Figure A1. Log-Loss for Different Hyperparameter Values

This figure plots the average log-loss, $\frac{1}{N}\sum_i^N [y_i \log p_i + (1 - y_i)\log(1 - p_i)]$ where p_i is the predicted probability of $y_i = 1$, among the cross-validation sets. Log-losses are computed for different learning rates and number of estimators while fixing other hyperparameters as default values.

The final hyperparameter values are marked in bold in Table A3. A more intensive search and optimization is not performed because it is not the main forcus this paper and because the bottleneck of performance is likely the small sample size. The final model has an average accuracy of 0.63, precision of 0.61, and AUC of 0.67 using a 5-fold cross-validation. I use a more conservative predicted probability cutoff of 0.65 to get the final result reported in the first column of Table 3.

As a reference, Figure A2 shows the mean decrease accuracy (MDA) for each feature. MDA is defined roughly as the loss in prediction accuracy, relative to the full model, when a given feature is randomly permuted. This measure is accepted to be most relevant for out-of-sample predictions because the computation of MDA naturally involves out-of-sample predictions. Figure A2 shows that the most relevant features are size, 5-minute volume, and minutes from open. A value of 0.03 means that permuting a given feature decreases the prediction accuracy by 3 percentage points. The final model excludes features that have a negative MDA value.

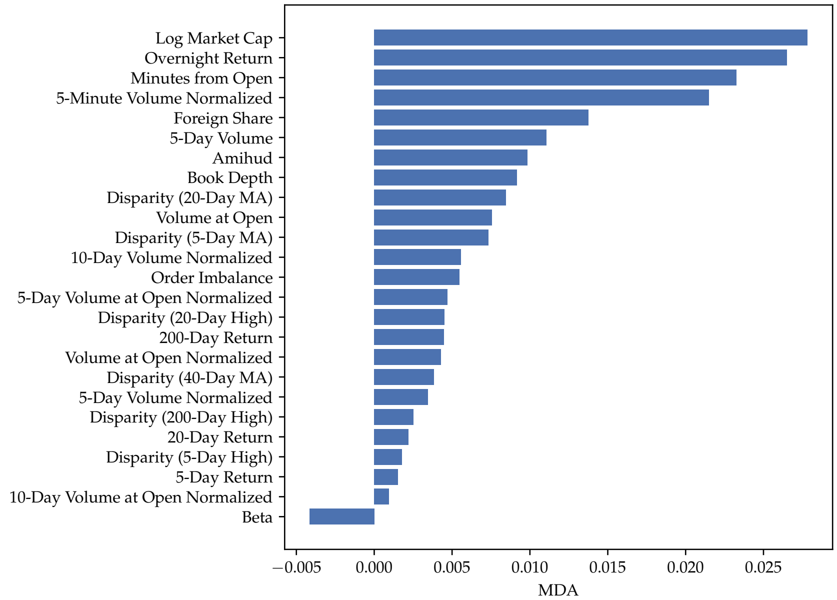

Figure A2. Mean Decrease Accuracy of Features

This figure illustrates the MDA of the features used in the prediction exercise. 10 test samples were created using 0.2 of the data for each test sample. For each train-test set, a model was trained with one randomly permuted column and its performance was compared to the full model on the same train-test set. The average differences on the 10 sets are recorded and plotted.

Table A3. Grid Search for Hyperparameters

Hyperparameter Name	Definition	Grid
Learning Rate	Step size shrinkage used in update to prevent overfitting.	[**0.01**, 0.05, 0.1, 0.15, 0.2, 0.3]
Number of Estimators (+)	Number of weak learners (trees).	[100, 200, 300, 400, **500**]
Maximum Depth (−)	Maximum depth of a tree.	[3, 4, **5**, 6, 7]
Minimum Loss Reduction (+)	Minimum loss reduction required to make a further partition on a leaf node of the tree.	[0, 0.25, **1**, 10]
Subsample Ratio (−)	Subsample ratio of the training instances at each tree growth.	[0.5, 0.7, **0.9**]
Column Subsample Ratio (−)	The subsample ratio of columns (features) when constructing each tree.	[0.3, **0.5**, 0.8]

This table lists the hyperparameters of XGBoost. Definitions are taken from the documentation page. Sign next to the hyperparameter name indicates the direction of the hyperparameter value that leads to the classifier's variance reduction.

Appendix C Additional Figures and Tables

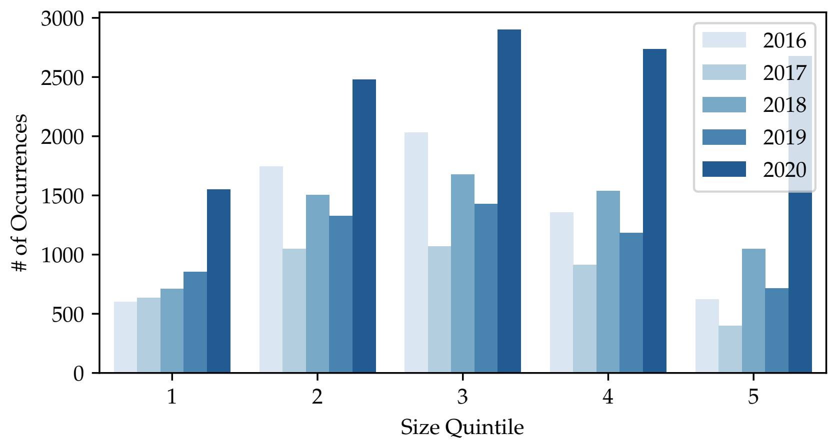

Figure A3. Number of VIs by Size

This figure illustrates the number of annual VI occurrences by size quintile. Bins for size were made based on market capitalization at the beginning of each month. VIs include both static and dynamic VIs.

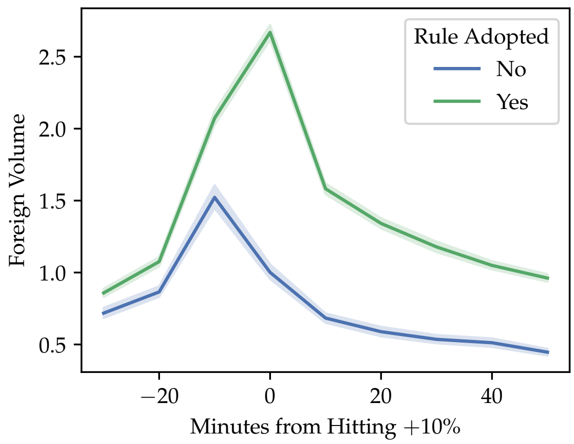

(a) Volume around Static-Up VI

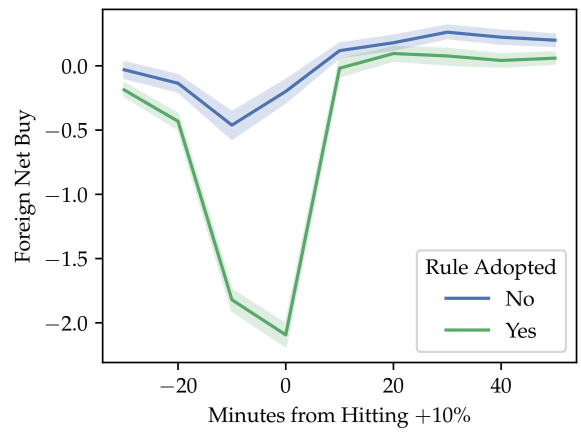

(b) Net Buy around Static-Up VI

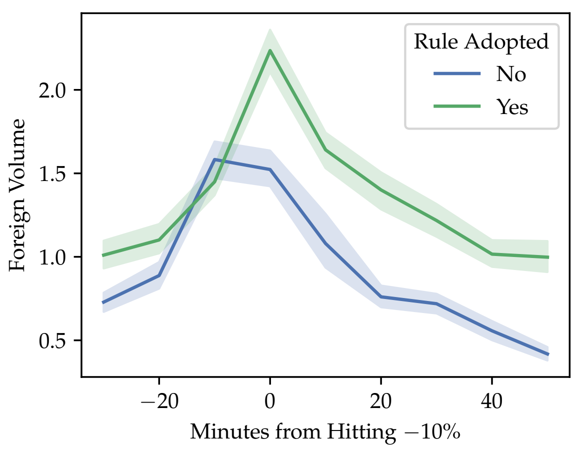

(c) Volume around Static-Down VI

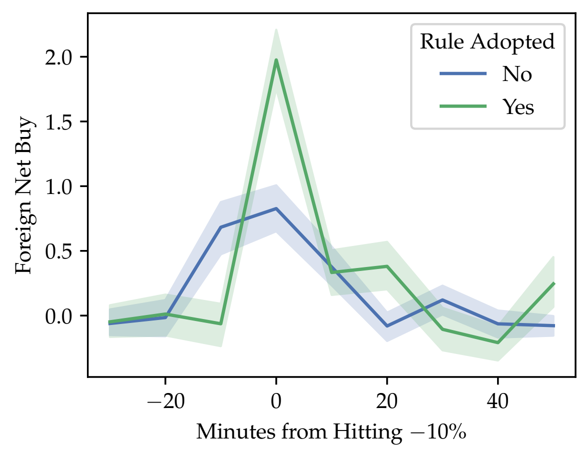

(d) Net Buy around Static-Down VI

Figure A4. Foreign Volume and Net Buy around Static VIs

This figure plots the average foreign institutions' volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

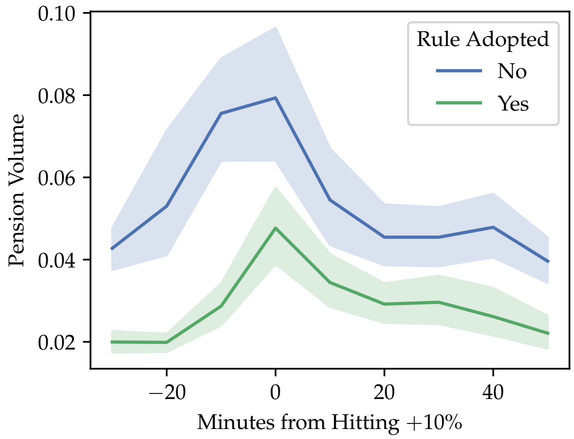

(a) Volume around Static-Up VI

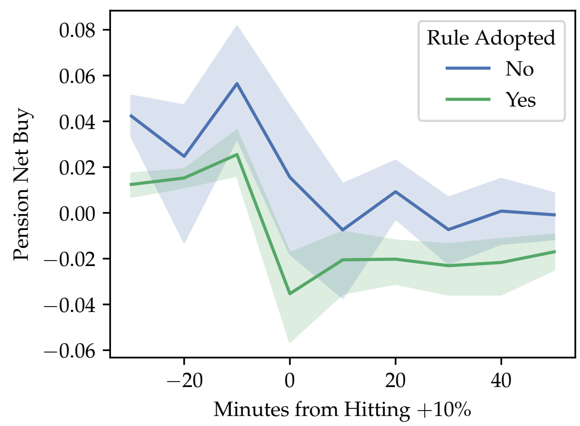

(b) Net Buy around Static-Up VI

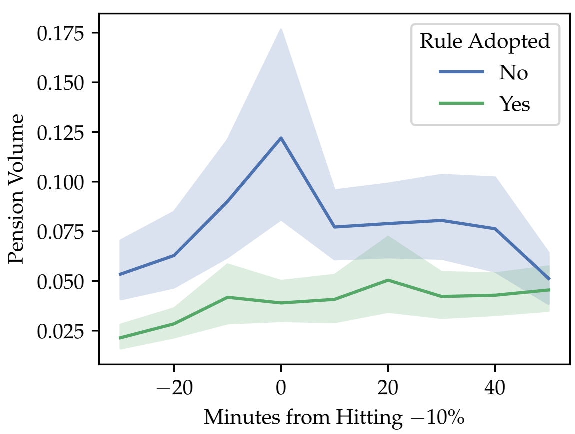

(c) Volume around Static-Down VI

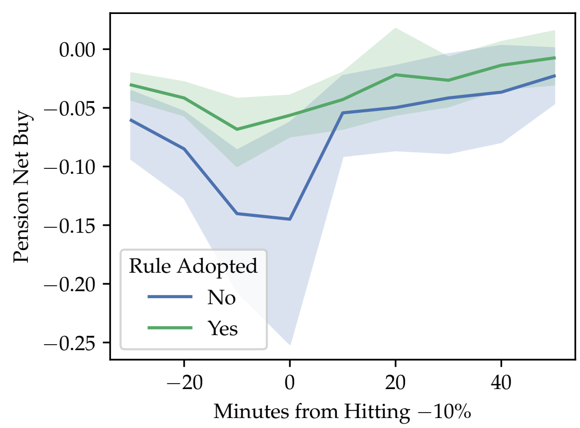

(d) Net Buy around Static-Down VI

Figure A5. Pension Volume and Net Buy around Static VIs

This figure plots pension funds' average volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

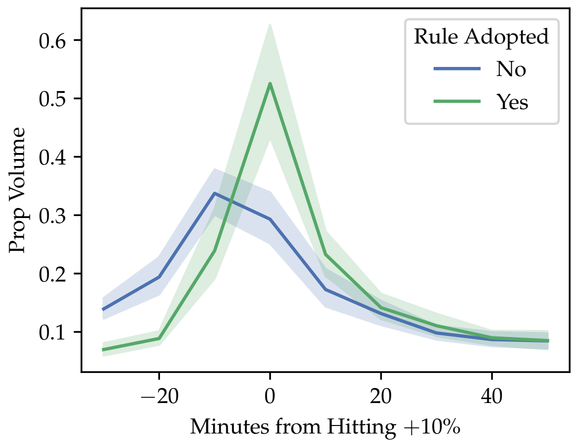

(a) Volume around Static-Up VI

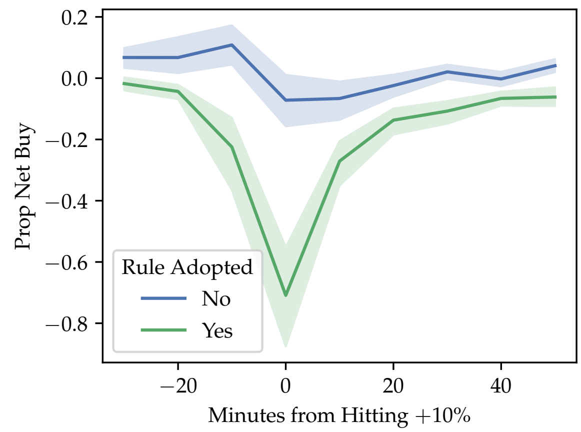

(b) Net Buy around Static-Up VI

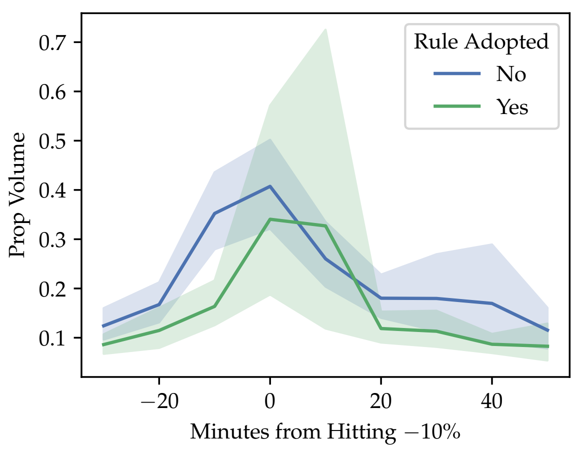

(c) Volume around Static-Down VI

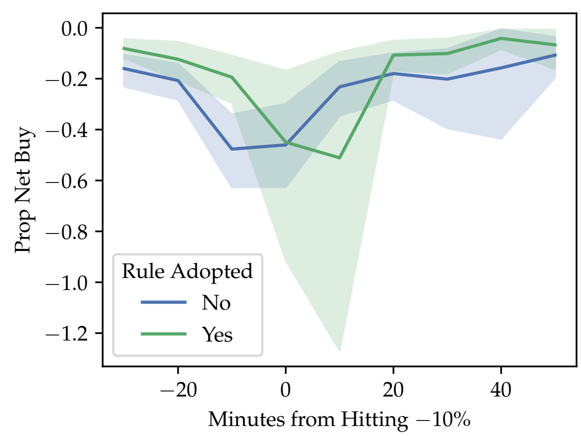

(d) Net Buy around Static-Down VI

Figure A6. Prop Volume and Net Buy around Static VIs

This figure plots proprietary traders' average volume and net buy during 10-minute bins around VIs and pseudo-pauses. Trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. Afterwards, they are expressed in basis points. The averages are taken for the pre-rule period (blue line) and post-rule period (green line). 95% confidence bands are constructed by bootstrapping.

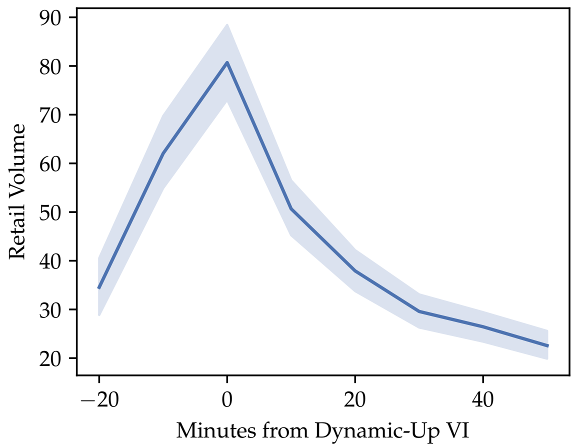

(a) Retail Volume

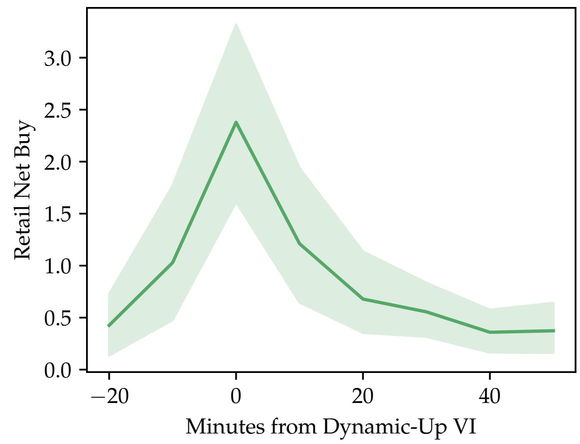

(b) Retail Net Buy

Figure A7. Retail Volume and Net Buy around Dynamic-Up VIs

This figure plots the average retail volume and net buy during 10-minute bins around dynamic-up VIs. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

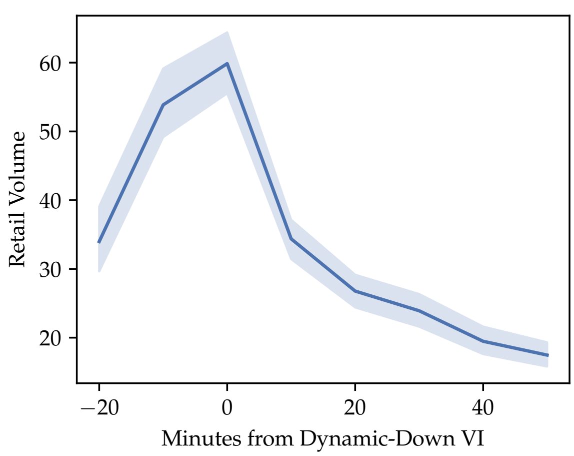

(a) Retail Volume

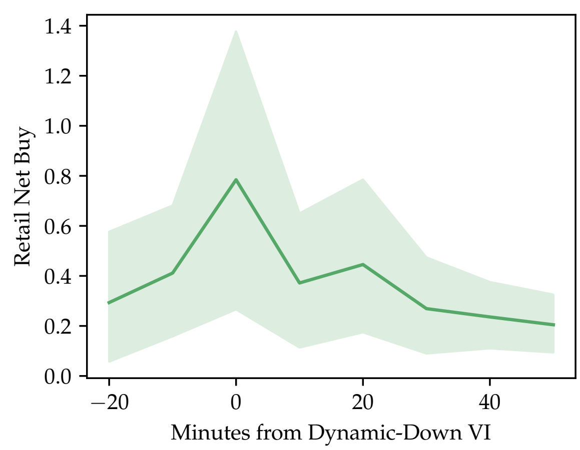

(b) Retail Net Buy

Figure A8. Retail Volume and Net Buy around Dynamic-Down VIs

This figure plots the average retail volume and net buy during 10-minute bins around dynamic-down VIs. Retail trading volume and net buy in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

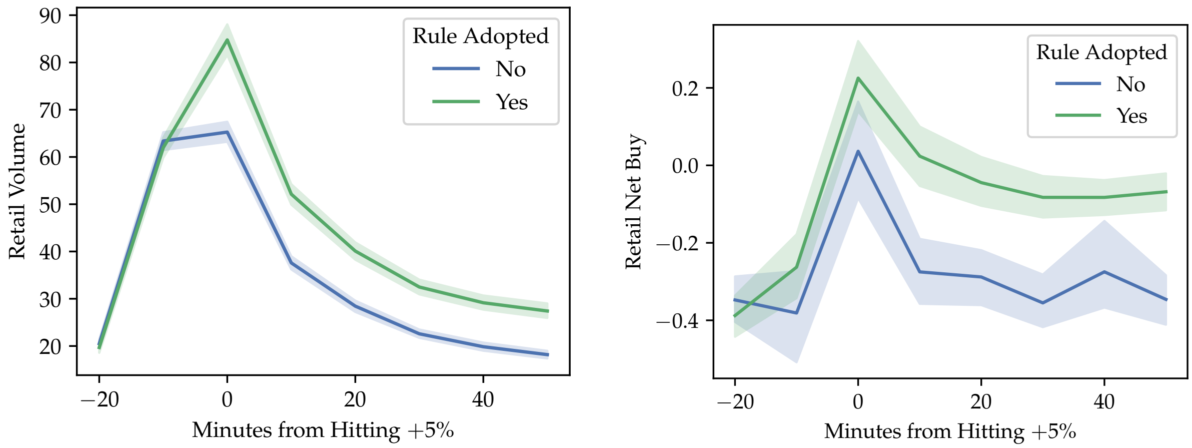

(a) Retail Volume

(b) Retail Net Buy

Figure A9. Retail Volume and Net Buy around +5% Breaches

This figure plots the average retail volume and net buy during 10-minute bins around +5% breaches. Retail trading volume in terms of shares are aggregated within a given stock-bin and normalized by number of shares outstanding. 95% confidence bands are constructed by bootstrapping.

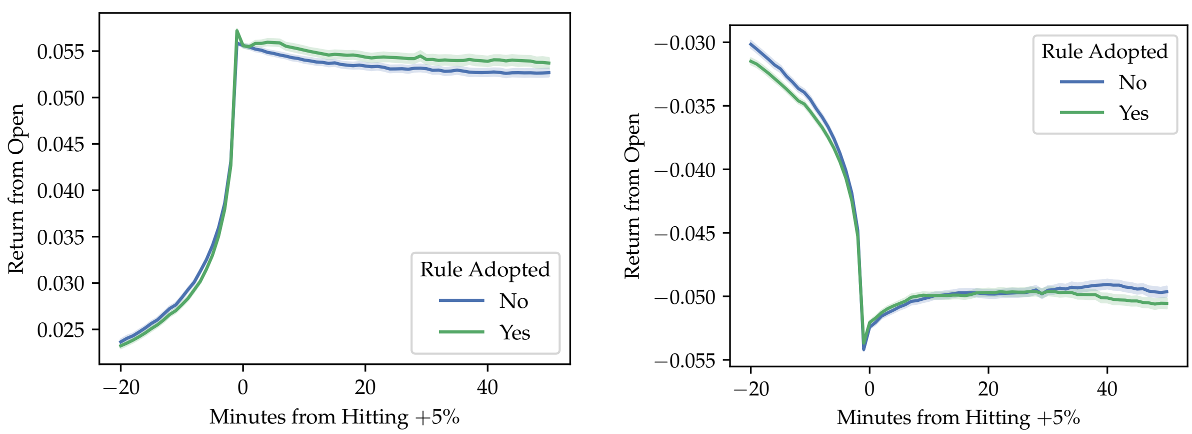

(a) +5% Breach

(b) −5% Breach

Figure A10. Prices around ±5% Breaches

This figure plots average prices around ±5% breaches, in terms of return from open. Return from open is computed using the closing price of each 1-minute interval according to Section 2.2. 95% confidence intervals are constructed by bootstrapping.

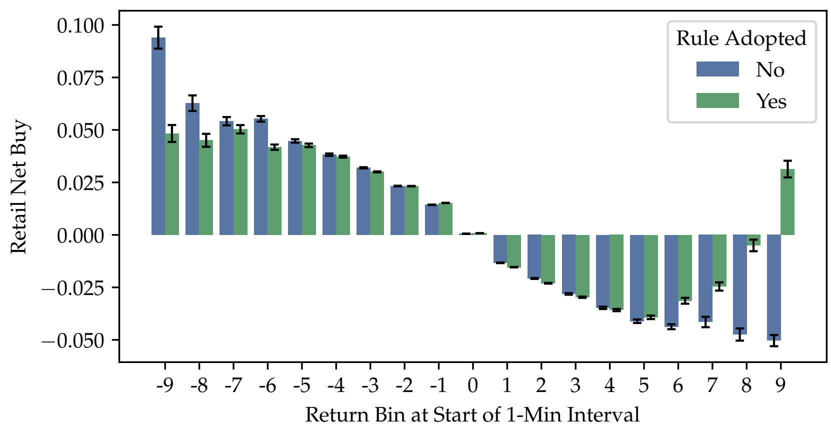

Figure A11. Retail Net Buy by Return Bin

This figure plots the average 1-minute net buy conditional on the return from open at the beginning of the 1-minute interval. The averages are computed for both pre-rule (blue bars) and post-rule (green bars) periods. Bars at $x = b$ refers to the group of intervals whose price, in terms of return from open, at the beginning fall between [b%, $(b+1)$%]. 95% confidence bands are constructed by bootstrapping.

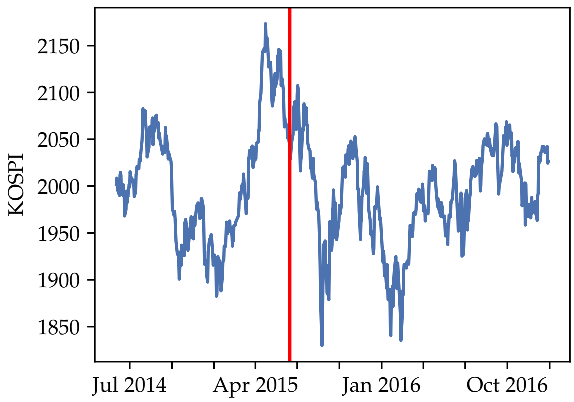

(a) KOSPI

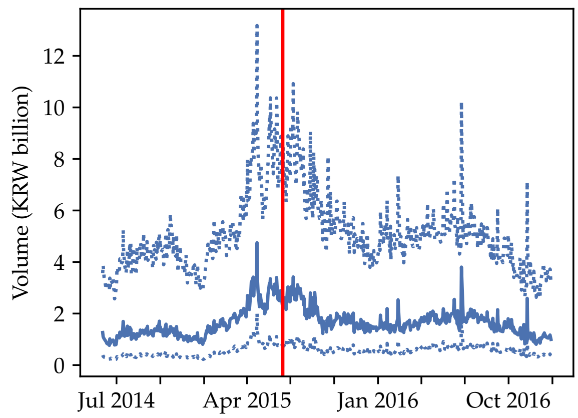

(b) Volume Quartiles

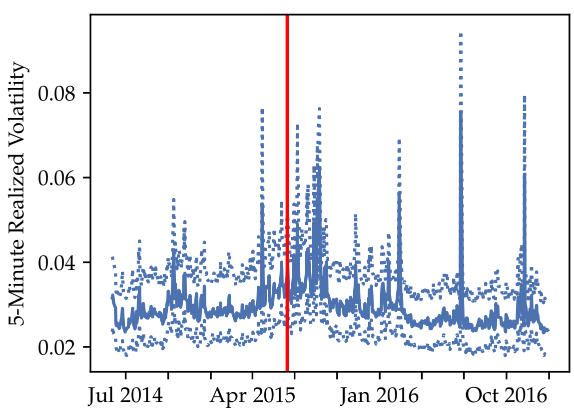

(c) Realized Volatility Quartiles

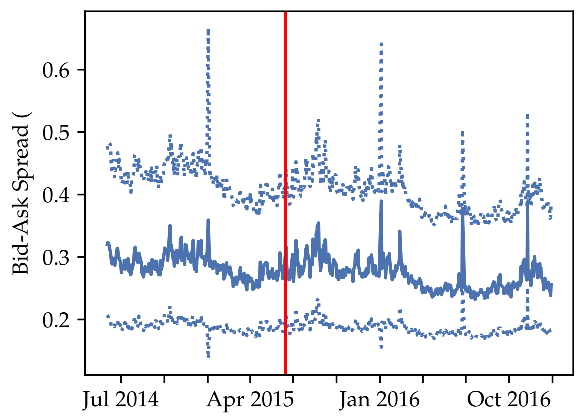

(d) Bid-Ask Spread Quartiles

Figure A12. Market Conditions around Rule Adoption

This figure plots the overall market conditions around the rule adoption. Panel (a) shows the level of KOSPI. Panel (b)–(d) plots the 25th, 50th, and 75th percentiles of volume, 5-minute realized volatility, and daily average bid-ask spreads in the cross-section of stocks.