

No. 571

Nikolaus Hautsch and Akos Horvath

How Effective are Trading Pauses?

The CFS Working Paper Series

presents ongoing research on selected topics in the fields of money, banking and finance. The papers are circulated to encourage discussion and comment. Any opinions expressed in CFS Working Papers are those of the author(s) and not of the CFS.

The Center for Financial Studies, located in Goethe University Frankfurt's House of Finance, conducts independent and internationally oriented research in important areas of Finance. It serves as a forum for dialogue between academia, policy-making institutions and the financial industry. It offers a platform for top-level fundamental research as well as applied research relevant for the financial sector in Europe. CFS is funded by the non-profit-organization Gesellschaft für Kapitalmarktforschung e.V. (GfK). Established in 1967 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia. GfK members comprise major players in Germany's financial industry. The funding institutions do not give prior review to CFS publications, nor do they necessarily share the views expressed therein.

How Effective are Trading Pauses?

Nikolaus Hautsch[†] Akos Horvath[‡]

First Draft: February 2015

This Draft: April 2017

Abstract

Exploiting NASDAQ order book data and difference-in-differences methodology, we identify the distinct effects of trading pause mechanisms introduced on U.S. stock exchanges after May 2010. We show that the mere existence of such a regulation constitutes a safeguard which makes market participants behave differently in anticipation of a pause. Pauses tend to break local price trends, make liquidity suppliers revise positions, and enhance price discovery. In contrast, pauses do not have a “cool off” effect on markets, but rather accelerate volatility and bid-ask spreads. This implies a regulatory trade-off between the protective role of trading pauses and their adverse effects on market quality.

Keywords: trading pause, “magnet effect”, price discovery, volatility, liquidity

JEL Classification: G10, G14, G18.

[†]Faculty of Business, Economics and Statistics, University of Vienna, Oskar–Morgenstern Platz 1, A-1090 Vienna, Austria. E-Mail: nikolaus.hautsch@univie.ac.at

[‡]Faculty of Business, Economics and Statistics, University of Vienna, Oskar–Morgenstern Platz 1, A-1090 Vienna, Austria. E-Mail: akos.horvath@univie.ac.at

1 Introduction

The statutory objectives of financial market regulation are to maintain market confidence and stability, protect consumers, and reduce financial crime (U.S. Financial Services and Markets Act, 2000). To promote a well-functioning market, control measures are imposed to regulate trading behavior and market activity. In this context, price limits and trading halts are implemented to prevent market participants from abrupt price changes. Given the extent of algorithmic (high-frequency) trading on nowadays markets, trading halts serve as important safeguards to protect market stability and to limit the risk of algorithm-induced overreactions or break-downs in liquidity supply and demand.

In response to the flash crash in May 2010, when the Dow Jones index realized the fastest price decline in history, trading pauses (also known as “volatility circuit breakers”) have been introduced on U.S. stock exchanges. A trading pause is a five-minute trading interruption which is automatically triggered for any listed stock, should its price change too rapidly. During this period, order execution is suspended, while order submissions and cancellations are still possible. The purpose of the temporary interruption of trading activity is to provide market participants time to process information and revise open positions. According to the U.S. Securities and Exchange Commission (SEC), see SEC (2012), this regulatory measure is expected to curb excessive stock market volatility and to serve as an additional layer of investor protection.

Market microstructure theory is generally undecided about the effects of circuit breakers. While Subrahmanyam (1994) argues that artificial constraints on price movements may have the undesirable effect of exacerbating the situation and self-fulfilling a market failure, Kyle (1988) argues that trading suspensions have a pacifying effect and help resolving market tensions. Existing empirical research, however, has primarily focused on the effects of trading *halts* (Lee et al., 1994; Corwin and Lipson, 2000; Christie et al., 2002), a former control mechanism, different from trading pauses and price limits (Harris, 1998; Cho et al., 2003). Trading halts suspend all order book activity and are activated at the discretion of the exchange for a longer time period (sometimes even for hours). In contrast, trading *pauses* only suspend order execution, are automatically triggered by a strict and publicly known rule, and last (according

to the current regulation) for exactly five minutes. Their effects on market quality, particularly in light of nowadays automatized trading, however, are widely unexplored.¹.

Exploiting a unique database, this paper provides novel insights into the effects and effectiveness of trading pauses. We aim at addressing the following research questions: (i) Do trading pauses contribute to a market “cool-off” by reducing volatility? (ii) How do trading pauses affect market liquidity? (iii) Do trading pauses facilitate trend reversals and price discovery, and (how) do market participants use these trading interruptions for strategic re-positioning?

To address these research questions, we collect the exact date and time of all trading pauses on the NASDAQ stock exchange from June 2010 (when the regulation came into force) to June 2014. We utilize twenty levels of the limit order book, which is updated at message-level frequency, enabling us to monitor market activity on a microsecond basis. The major empirical challenge, however, is to isolate the effect of trading pauses from the effect of the underlying price movements triggering the pauses. In order to identify what would have happened without a trading pause, we construct a control set of price change periods for the corresponding stocks during the pre-regulation time period, when trading pauses have not been implemented yet. We show that the pre-regulation and post-regulation periods are similar in terms of volatility and market liquidity, which makes the market scenarios around extreme price movements well comparable. Performing difference-in-differences analysis and applying high-frequency econometrics, we quantify the effect of trading interruptions on short-term volatility, bid-ask spreads, liquidity demand, and liquidity supply.

An advantage of NASDAQ limit order book data is that it provides information on quoting behavior *during* trading pauses. This gives us the unique opportunity to analyze how market participants position themselves during pauses, in particular, when the end of the trading interruption, and thus the resumption of trading, is approaching. Since there is no order-matching (i.e. execution) during the trading pause, we synthetically clear the market, and compute the implied mid-quotes from the (hypothetical) intra-pause limit order book. This way, we examine the informativeness of limit orders and the price formation during trading pauses,

¹The fact that the SEC extended the initial pilot period for the evaluation of trading pauses, indicates that a deeper understanding of the impact of these regulatory measures is still needed, see SEC (2016)

and assess to what extent intra-pause mid-quotes predict price levels *after* the pause. Such an analysis gives novel insights into price discovery as well as the speed of price convergence during and after pauses.

Our empirical results show that trading pauses generally fulfill an important role as circuit breakers of local price trends and serve as safeguards in volatile market periods. We provide evidence that the mere presence of such a trading pause regulation changes the trading and quoting behavior of market participants in periods of strongly moving prices. Our findings indicate that the existence of a trading pause mechanism creates a layer of protection for market participants who are willing to trade against the price movement and initiate a trend reversal. In particular, we find that such counter-acting forces in liquidity supply become active already prior to trading pauses. This is only true, however, if a trading pause regulation is in force, and thus market participants obviously know that they are protected from adverse selection under extreme market conditions.

Analyzing market participants' quoting behavior during trading pauses shows that the counter-acting shifts in liquidity supply against the direction of the underlying price trend become stronger after the interruption of trading. Hence, liquidity suppliers efficiently use trading pauses to moderate price movements and facilitate mid-term price stabilization. Accordingly, limit orders posted during trading pauses are longer valid than limit orders posted outside trading pauses, which makes implied mid-quotes informative for (temporary) stable price levels after the trading interruption. This re-positioning of liquidity supply converges quickly, which implies price discovery at a new level already after the beginning of the pause. Our results therefore suggest that it is rather the existence of a pause mechanisms than its actual length, which makes a trading pause functioning as a price trend breaker and safeguard against adverse selection. Supported by the observed intra-pause order arrival patterns, and based on the theoretical arguments by Glosten (1994), we argue that this protective function of trading pauses encourages particularly informed traders to become counter-active and to push prices towards (new) equilibrium levels.

We further show that a trading pause regulation also has adverse effects on market quality. While it serves as a safeguard, making market participants to behave differently already in anticipation of a trading suspension, it also creates excess volatility. In the spirit of the “magnet

effect”, described by Subrahmanyam (1994), the possibility of the occurrence of a trading pause significantly increases volatility and trading volume prior to the pause and up to thirty minutes thereafter. Therefore, our results indicate that trading pauses do *not* serve as volatility circuit breakers, and do not help the market “cool off” in terms of reduced price uncertainty. We rather observe that the pause-induced break of local price trends is accompanied by increased volatility, not only after the trading break, but also prior to it. Moreover, this effect goes hand in hand with a widening of bid-ask spreads and thus an increase in transaction costs.

These findings are important for regulators, as they show that there is no “free lunch” in financial market regulation. Trading pauses artificially interrupt the flow of trading. This artificial disruption obviously causes temporary disturbances which materialize in higher volatility and transaction costs. In this sense, our results confirm the ambiguity in the theoretical market microstructure literature, supporting arguments both in favor of and against a trading pause regulation. The results show that such a regulatory measure has to be implemented with caution, owing to the natural trade-off between its function as a safeguard and its adverse temporary effects on market quality.

This paper is structured as follows. In Section 2, we review the most important economic arguments and empirical findings, and discuss the regulatory perspective behind trading pauses. Section 3 presents the data and an explorative analysis of volatility and liquidity around trading pauses. In Section 4, we present the estimates of the distinct effects of trading pauses on volatility, order book liquidity and market asymmetry. In Section 5, we study quoting behavior and price discovery *during* trading pauses. Finally, Section 6 concludes.

2 Theory and Practice of Circuit Breakers

2.1 Literature

Market microstructure theory is undecided about the potential costs and benefits of circuit breakers. Brady (1988) and Kyle (1988) argue that trading halts reduce volatility, resolve order imbalances, limit credit risk, and protect liquidity providers by allowing market participants to process information, to revise positions, and to collect margins to cover incurred losses. In their

view, this additional layer of security has a positive effect on market liquidity and facilitates arbitrage by synchronizing market forces.

These arguments are consistent with the model by Greenwald and Stein (1991), suggesting that trading halts enhance price discovery by trading off the immediacy of a continuous market against a more complete information transmission between market participants. Likewise, Madhavan (1992) claims that continuous trading may not be practical during periods of severe information asymmetry, and recommends a temporary switch to periodic trading, so that sufficient liquidity can accumulate on both sides of the market.

On the other hand, according to Fama (1989), trading halts do not reduce but merely postpone volatility, which spills over to the period after the forced interruption of stock price movements has been lifted. In the model of Subrahmanyam (1994), the very presence of trading interruption rules amplifies shocks to the stock price by causing traders to cancel positions and to withdraw from the market. This self-exciting mechanism is called the “magnet effect”, as it increases volatility and pushes the stock price towards the threshold where the trading interruption is triggered.

Harris (1998) takes an intermediate position, arguing that volatility can originate from both fundamental (informed) and transitory (uninformed) trading. He claims that trading halts may help to resolve order imbalances and reduce volatility if uncertainty is driven by uninformed noise. At the same time, price constraints have a detrimental effect on price discovery when volatility is induced by informed trading, and thus circuit breakers might unnecessarily inhibit the spread of information.

Empirical research tried to shed light on the ambiguous theoretical predictions, mainly focusing on trading *halts*. Based on data from the NYSE, Lee et al. (1994) detect increased post-halt trading volume, which is confirmed by Corwin and Lipson (2000) and Christie et al. (2002), who find higher bid-ask spreads and volatility for at least two hours after trading halts. Another aspect of trading halts is addressed by Bhattacharya and Spiegel (1998), who investigate post-halt stock return patterns on the NYSE in the period 1974-1988. They find evidence for the theoretical prediction that trading halts triggered by news arrivals are followed by return continuation, whereas trading halts caused by market imbalances are followed by return reversal.

Abad and Pascual (2010) examine changes in market liquidity around “volatility call auctions”, which were introduced in 2001 on the Spanish Stock Exchange and are triggered automatically based on price movements. Taking 5-minute sampling windows in the period 2001-2003, they find that trading volume and mid-quote price ranges significantly increase during these periods, slowly returning to their initial levels over the next two hours.

Only a few papers studied more recent automatized control mechanisms. Zimmermann (2013) investigate volatility interruptions on the Xetra stock exchange, which are similar to NASDAQ trading pauses. They find that price uncertainty prior to the interruption is significantly reduced by the break and interpret this as a sign of improved price discovery. Brogaard and Roshak (2015) investigate the effects of NASDAQ trading pauses on the frequency and severity of extreme price movements, comparing pre-regulation and post-regulation stock returns on an aggregated basis. They find that trading pauses reduce the size of extreme price movements, but induce price under-reactions.

Intra-day effects of trading pauses on volatility, liquidity and price discovery in nowadays high-frequency market environments, however, are still widely unexplored. We contribute to the empirical literature by making use of unique limit order book data.

2.2 Regulatory Rules

After the Flash Crash of May 6, 2010, the U.S. Commodity Futures Trading Commission (CFTC) and the U.S. Securities and Exchange Commission (SEC) assembled a joint advisory committee to find potential regulatory measures for improving stock market stability. Their report (CFTC and SEC, 2010) calls attention to the growing importance of high-frequency trading. It is further pointed out that (i) trading volume is not necessarily a reliable indicator of market liquidity, (ii) high-frequency trading strategies can quickly erode liquidity, and (iii) stock and derivatives markets are strongly interconnected with respect to liquidity. These factors together can lead to a sudden evaporation of liquidity from the market, which results in erratic price movements, causing losses to investors.

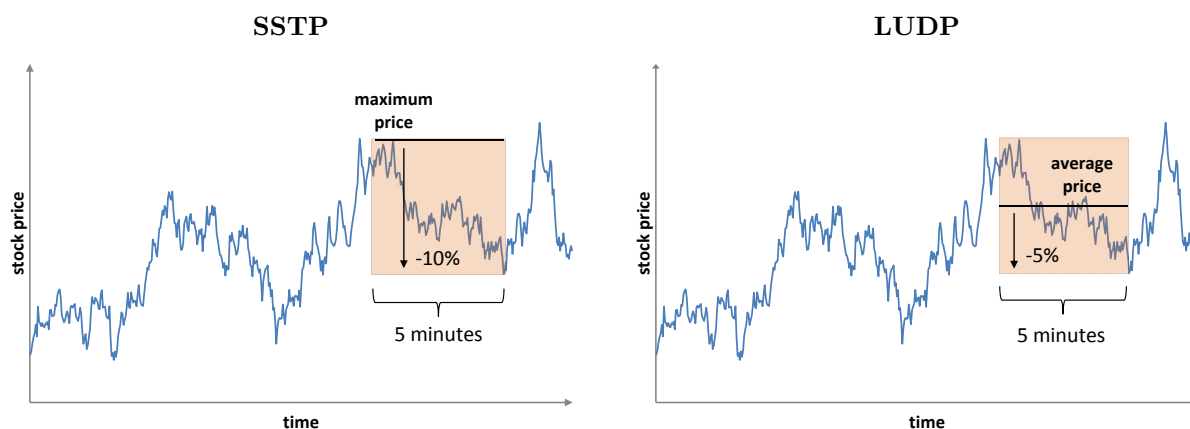
To avoid this undesirable outcome and prevent future market crashes, on June 10, 2010, the SEC, collaborating with the Financial Industry Regulatory Authority (FINRA), approved

the implementation of the so-called single-stock trading pauses (SSTPs). These are five-minute call auctions, which are individually and automatically triggered for any listed stock, should the stock price change too quickly. A crucial difference to existing trading halt mechanisms is that the latter are activated at the discretion of the individual stock exchange and do not allow for changes of the order book during the pause. In contrast, SSTPs simultaneously affect *all* U.S. exchanges and suspend order execution, but not order placement and quoting activity. According to the SEC (see SEC (2010)), they serve the purpose of “providing market participants more certainty, as to which trades will be broken, and allow them to better manage their risks”.

On April 5, 2011, the SEC approved certain amendments to the original pausing rule and introduced the so-called limit-up limit-down trading pauses (LUDPs), which gradually superseded SSTPs as of April 2013, see SEC (2011). The respective trigger rules are illustrated in Figure 1. In the case of SSTPs, trading is paused whenever the current stock price undercuts (or exceeds) the maximum (or minimum) price of the last five minutes by more than 10%. In contrast, LUDPs suspend trading if the current price leaves the $\pm 5\%$ corridor around the *average* price of the last five minutes.

Figure 1: Trading Pause Trigger Rules

This figure illustrates the trading pause rules by means of an example price path where the trading interruption is actually triggered. SSTP stands for single-stock trading pauses, coming into force in June 2010, and LUDP stands for limit-up limit-down trading pauses, coming into force in April 2013, and superseding the earlier stopping mechanism. Both trading pauses are automatically triggered based on the last five minutes of stock price history, which is portrayed by the shaded rectangles.



3 Data and Descriptive Statistics

3.1 NASDAQ Trading Pauses

Using the information acquired from `NASDAQTrader.com`, we determine the exact date and time of every SSTP and LUDP trading pause that occurred on NASDAQ since their implementation in June 1, 2010 to June 30, 2014. Overall 1513 trading pauses occurred in this time period, which are characterized in Table 1. We find that trading pauses are relatively equally distributed across stocks, as indicated by the number of affected stocks as well as metrics of concentration. Moreover, the LUDP rule results in more frequent trading pauses than the SSTP rule, which is likely due to the slower response of the average price, rendering the LUDP more sensitive to temporary price fluctuations.

Table 1: Occurrence of Trading Pauses

This table shows summary statistics on the occurrence of trading pauses in the time period from June 1, 2010 to June 30, 2014. SSTP stands for single-stock trading pauses, coming into force in June 2010, and LUDP stands for limit-up limit-down trading pauses, coming into force in April 2013, and superseding the earlier stopping mechanism. The frequency of trading pauses is calculated on a yearly basis (i.e. pauses per stock per year), whereas the concentration across individual stocks is characterized by the Herfindahl–Hirschman index (hhi) as well as the Gini coefficient (gini).

	total number of		pauses per stock (per year)				concentration	
	pauses	affected stocks	mean	stdev	min	max	hhi	gini
SSTP	400	233	0.57	1.21	0.33	9.67	0.011	0.333
LUDP	1113	242	4.60	10.20	1.00	107.00	0.024	0.654

Figure 2 shows the distribution of SSTPs and LUDPs across and within trading days. Even though there is no detectable pattern across the days of the week, it turns out that trading pauses occur more often in the morning hours. This is consistent with the well-known observation that intraday stock return volatility is highest during the morning. This connection between higher volatility and trading pause frequency is confirmed by Figure 3, which shows that the most volatile 20% of stocks trigger significantly more (about twice as many) trading pauses than other stocks with more stable prices.

Figure 2: Trading Pause Distribution by Occurrence Time

This figure shows average daily trading pause frequencies across and within trading days, computed based on the 1513 trading pauses occurring in the time period from June 1, 2010, to June 30, 2014. SSTP stands for single-stock trading pauses, coming into force in June 2010, and LUDP stands for limit-up limit-down trading pauses, coming into force in April 2013, and superseding the earlier stopping mechanism. For both the SSTP and the LUDP group, the equality of hourly (daily) mean frequencies is tested by analysis of variance (ANOVA). The resulting p-values are shown in the shaded rectangles.

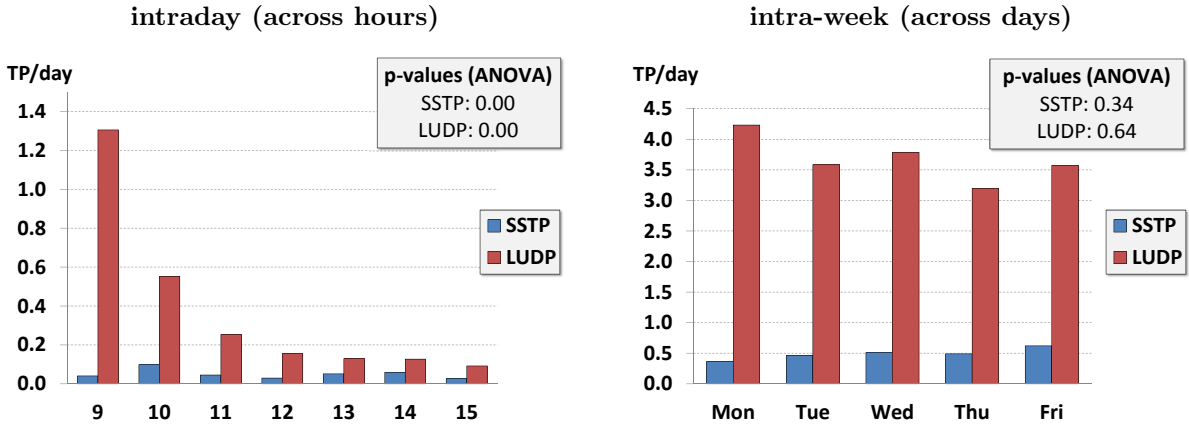
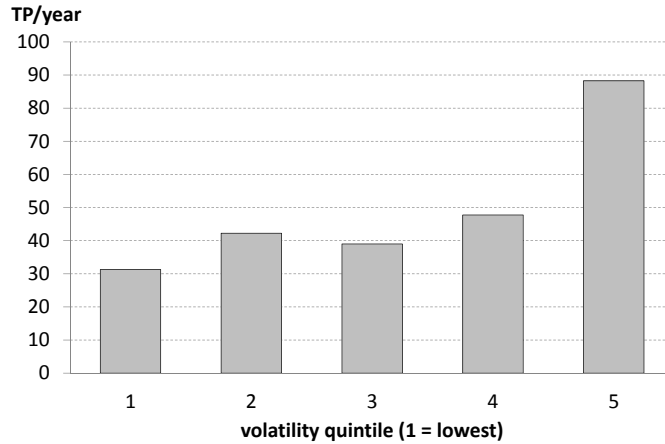


Figure 3: Trading Pause Distribution by Stock Volatility

This figure shows average yearly trading pause frequencies across realized volatility quintiles, computed based on the 1513 trading pauses occurring in the time period from June 1, 2010 to June 30, 2014. Realized volatility is computed for each paused stock from daily closing prices in this time period.




The information from `NASDAQTrader.com` is merged with underlying limit order book data from the platform `LOBSTER`², providing comprehensive message-level order book data for each stock listed on NASDAQ. Given the objectives of this paper, and the need to quantify the (high-

²For more details about the `LOBSTER` database, see <http://lobsterdata.com>.

frequency) evolution of volatility and liquidity around trading pauses with reliable statistical precision, we exclusively focus on trading pauses which are accompanied by a sufficiently high trading activity in the underlying stock in periods before and after the pause. We therefore only keep a trading pause in our sample if there are at least 300 limit order book messages in the hour before and after the pause. We also require that the standard deviation of inter-trade waiting times is less than 30 seconds. This way, we rule out that trading activity is clustered around only a few time points, and ensure that it is evenly spread across time.³ In case these two criteria are not fulfilled for the entire pre-pause and post-pause trading hours, but only for a shorter time period around the pause, we restrict the analysis to this shortened time period. As a consequence, for some trading pauses, shorter pre-pause and post-pause intervals are used. This way, we retain the maximum number of trading pauses in the sample, while ensuring that pre-pause and post-pause effects can be estimated with sufficient precision. Finally, we keep only those trading pauses for which at least a 5-minute pre-pause and a 20-minute post-pause trading period exist, and exclude those whose distance from the next trading pause is less than two hours, in order to avoid overlapping effects. Figure 4 illustrates how the pre-pause and post-pause intervals are defined.

Figure 4: Illustration of Pre-Pause and Post-Pause Time Intervals

This figure illustrates how the pre-pause and post-pause time periods are defined. The beginning and end of each five-minute trading pause is denoted by H0 and H1, respectively. Around each trading pause, we start with a two-hour interval (one hour before and one hour after), which is then shortened until the remaining pre-pause (from T0 to H0) and post-pause (from H1 to T1) time periods fulfill the used criteria for sufficient trading activity. This procedure is then performed for each 1513 trading pauses occurring in the time period from June 1, 2010 to June 30, 2014, enabling us to select those 195 trading pauses which have sufficiently long pre-pause and post-pause time periods for the objective of this analysis.



³The chosen criteria are obviously somewhat ad hoc. We find, however, that the sample constitution is robust to changes in these criteria, and does not seem to systematically bias our analysis. Filtering out certain trading pauses with insufficient trading activity is unavoidable, so as to ensure that the liquidity and volatility variables used in this study are numerically reliable and computationally tractable.

The final sample consists of 195 trading pauses, affecting 166 stocks. As shown in Figure 5, the pauses are spread quite evenly across the sample time period. Figure 6 reveals that trading pauses are relatively equally placed across price levels and market capitalization. Nevertheless, as a result of our pre-filtering, we do observe a slight tilt towards large-cap stocks.

Figure 5: Distribution of Trading Pauses by Calendar Date

This figure shows the distribution of the 195 selected trading pauses across calendar dates in the time period from June 1, 2010 to June 30, 2014. The selected trading pauses affect 166 stocks listed on NASDAQ. SSTP stands for single-stock trading pauses, coming into force in June 2010, and LUDP stands for limit-up limit-down trading pauses, coming into force in April 2013, and superseding the earlier stopping mechanism.

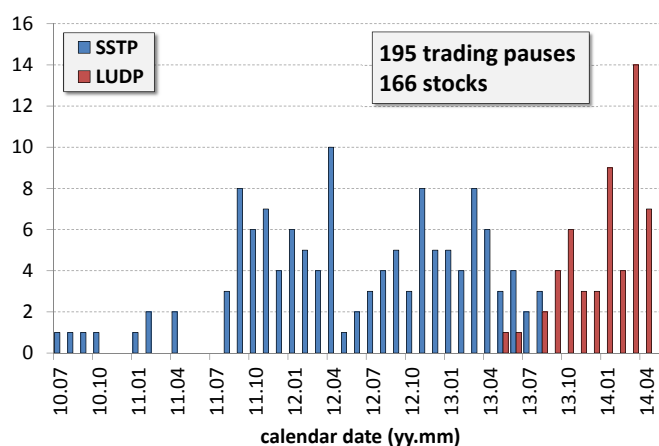
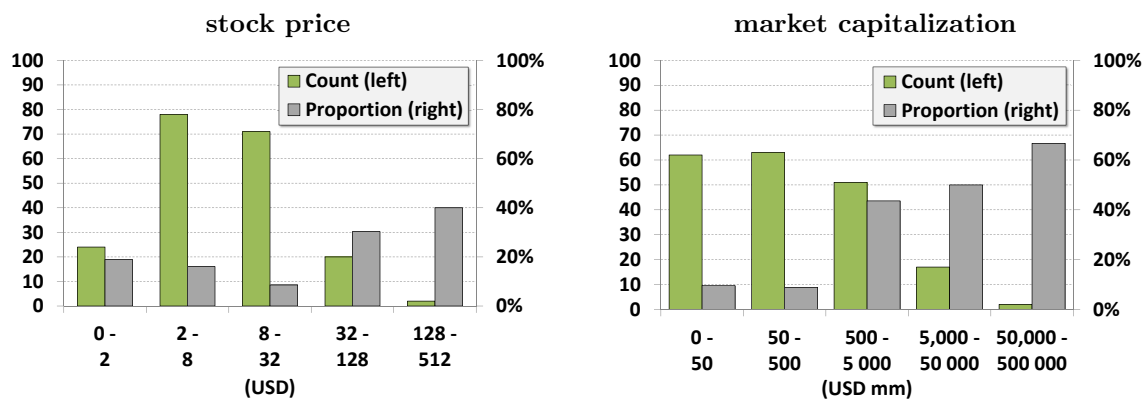


Figure 6: Distribution of Trading Pauses by Price Level and Market Capitalization

This figure shows the distribution of the 195 selected trading pauses across price levels and market capitalization. The counts are measured on the left axes, showing how many sample trading pauses fall into the respective bin. The relative proportion is measured on the right axes, showing the percentage of NASDAQ stocks affected by trading pauses in the respective bin.



3.2 Volatility and Liquidity Around Trading Pauses

We start our empirical investigation by an exploratory analysis of volatility and liquidity measures around trading pauses. To gain insights into the evolution of these quantities at a high-frequency level, we make use of the available limit order book data, and compute different volatility and liquidity statistics over 3-minute and 1-minute bins, respectively, in the two-hour interval around trading pauses defined above⁴.

To efficiently estimate local return (co)volatility over short-term intervals, we apply the methodology developed by Reiss (2011) and Bibinger et al. (2014), proposing state-of-the-art estimators to efficiently back out variances and covariances from noisy high-frequency data. Here, we briefly illustrate the main idea behind the estimator, using a simplified (univariate) setting of observed prices which are contaminated by i.i.d. noise. Observed log prices are assumed to be given by $y_i = x_i + \varepsilon_i$, $i \in \{1, \dots, n\}$, where x_i corresponds to the “efficient” (fundamental) price following a (continuous) martingale process $dX_t = \sigma(t)dB_t$, with Brownian motion B_t and spot volatility $\sigma(t)$. The i.i.d. error term $\varepsilon_i \sim N(0, \eta^2)$ is associated with market microstructure noise with variance η^2 . The estimator rests upon the idea of splitting the (normalized) period of interest $[0, 1]$ (in the given context, a 3-minute or 5-minute bin, respectively) into local intervals of length h , $[kh, (k+1)h]$, $k \in \{0, \dots, h^{-1} - 1\}$, and considering the (interval-specific) volatility process as being locally constant, that is $\sigma(t) = \sigma_k$ for $t \in [kh, (k+1)h]$. Estimators for the block-wise variances σ_k^2 can then be constructed based on local spectral statistics. Define the block-wise (spectral) statistic s_{jk} , computed based on a given spectral frequency $j \geq 1$, as

$$s_{jk} = \pi j h^{-1} \int_{kh}^{(k+1)h} \varphi_{jk}(t) dy(t), \quad j \geq 1, \quad k = 0, \dots, h^{-1} - 1,$$

where $\varphi_{jk}(t)$ denotes the orthogonal sine function with (spectral) frequency j , given by

$$\varphi_{jk}(t) = \frac{\sqrt{2}}{h} \cos(j\pi h^{-1}(t - kh)) \mathbb{1}_{[kh, (k+1)h)}(t), \quad j \geq 1.$$

⁴For the sake of robustness, we also repeated the empirical analysis using longer, 5-minute and 3-minute bins, respectively. The results are qualitatively and quantitatively similar.

Reiss (2011) shows that s_{jk} is independent for all j and k , and $s_{jk} \sim N(0, c_{jk})$ with variance $c_{jk} = \sigma_k^2 + \pi^2 j^2 h^{-2} \eta^2 / n$. This result implies that for each frequency j , an unbiased estimator σ_k can be constructed based on the bias-corrected empirical variance of s_{jk} , given by $s_{jk}^2 - \frac{j^2 \pi^2 \hat{\eta}^2}{h^2 n}$. Taking a weighted sum of these (moment) estimators for all spectral frequencies j and all blocks $k \in \{0, \dots, h^{-1} - 1\}$ yields a convex combination of bias-corrected empirical variances of s_{jk} , which serves as an estimator for the integrated variance $\sigma^2 := \int_0^1 \sigma^2(s) ds$ over the (normalized) interval of interest,

$$\hat{\sigma}^2 := \sum_{k=0}^{h^{-1}-1} h \sum_{j=1}^{\infty} w_{jk} \left(s_{jk}^2 - \frac{j^2 \pi^2 \hat{\eta}^2}{h^2 n} \right), \quad j = 1, \dots, nh - 1.$$

As shown by Reiss (2011), an optimal choice of the weights $w_{jk} = c_{jk}^{-2} (\sum_{j=1}^{\infty} I_{jk}^{-2})^{-1}$ minimizes the estimator's variance. If the block length shrinks sufficiently fast for $n \rightarrow \infty$, then the estimator $\hat{\sigma}^2$ is consistent and asymptotically efficient. Bibinger et al. (2014) propose a multivariate version of this estimator and discuss implementation details, such as the choice of the block length h , how many spectral functions j to include and how to (pre-)estimate the market microstructure noise variance η^2 .

We utilize a univariate version of the estimator based on mid-quote returns. Due to the high number of order book updates for liquid NASDAQ stocks, we are able to employ this estimator over 3-minute bins. Figure 17 in the Appendix illustrates the cross-sectional variation of high-frequency volatility around trading pauses, showing the evolution of the across-pause median and different quantiles across time. We observe that the median volatility increases about tenfold in the vicinity of trading pauses, momentarily reaching an annual level of 600% at the peak. Moreover, high-frequency volatility seems to increase further *after* the trading pause, in line with the theoretical arguments of Fama (1989) and Subrahmanyam (1994).

Implementing a bivariate version of the spectral estimator proposed by Bibinger et al. (2014), we study high-frequency correlations between paused stocks and the market as well as the corresponding industrial sector. As a proxy for the market, we use the exchange-traded fund ONEQ, tracking the NASDAQ Composite Index. As a proxy for the industrial sector, we construct an equal-weighted portfolio of the actively traded stocks with the same first three NAICS digits as

the corresponding paused stock. Using the market index and the industry portfolios, we construct covariance estimates over each 3-minute bin. The bottom of Figure 17 in the Appendix illustrates the cross-sectional variation of the resulting stock-market and stock-industry correlations around the trading pause. The observed correlation patterns indicate declining dependencies between paused stocks and other (non-paused) constituents of the market *after* the trading interruption.

Figure 18 in the Appendix shows the cross-sectional variation of (log) bid-ask spreads and (standardized) order book depth, calculated as the average of all order book updates for a series of 1-minute bins around trading pauses. We also illustrate the 1-minute (log) cumulative trading volume.⁵ Bid-ask spreads start to widen before the trading pause, reach a peak right at the trading interruption and gradually return to normal levels over the course of the following hour. Therefore, trading costs significantly increase prior to the pause, but gradually decline thereafter.

Market depth, aggregated over the first two limit order book levels and standardized by the monthly average for the respective stock, drops by approximately 25% after the trading pause is triggered, and remains reduced for at least one hour afterwards. We also observe that the median market depth is generally lower on days when trading pauses occur, reaching only about 85% of the depth on days without a trading pause. While liquidity supply drops after the pause, we observe that liquidity demand, measured by the cumulative trading volume peaks prior to the pause, and stays elevated over the next hour. Such a pattern is likely due to a positive relationship between volatility and trading activity.

In order to analyze whether the evolution of market depth (i.e. liquidity supply) and trading volume (i.e. liquidity demand) occurs on specific sides of the market, we construct limit order book imbalance and market order flow measures which incorporate the direction of the price movement. In particular, we define the directional (spread-weighted) limit order book imbalance (IMBAL) measure as

$$\text{IMBAL} := \text{DIR} \times \frac{\sum_{l=1}^2 d(l)^b(p(1)^a - p(l)^b) - d(l)^a(p(l)^a - p(1)^b)}{\sum_{l=1}^2 d(l)^b(p(1)^a - p(l)^b) + d(l)^a(p(l)^a - p(1)^b)}, \quad (1)$$

⁵We take the logarithm of the bid-ask spread and the trading volume to reduce the cross-sectional heterogeneity in these variables, and ease the comparison across different stocks.

where $p(l)^b$ and $p(l)^a$ denote the bid and ask quotes at the l^{th} order book level, while $d(l)^b$ and $d(l)^a$ represent the corresponding order book depth. Variable DIR takes the value of -1 or $+1$, corresponding to the direction of the price movement triggering the trading pause. The order book imbalance measure is averaged over all limit order book updates in each 3-minute bin. Similarly, directional market order flow (OFLOW) is defined as

$$\text{OFLOW} := \text{DIR} \times \frac{Q(\text{buy}) - Q(\text{sell})}{Q(\text{buy}) + Q(\text{sell})}, \quad (2)$$

where $Q(\text{buy})$ and $Q(\text{sell})$ represent the cumulative volume of buyer and seller initiated market orders in each 3-minute bin.

By construction, both measures IMBAL and OFLOW are bounded between -1 and $+1$, which simplifies the comparison and aggregation across paused stocks with different order book depth. Positive values of these statistics indicate market pressure in the direction of the price movement, whereas negative values indicate market pressure in the opposite direction. From Figure 19 in the Appendix, however, we cannot identify strong market side-specific patterns in the direction of the price movement. There is a slight increase of order flow in the direction of the price change prior to the pause, indicating buying (selling) pressure for upward (downward) price movements. Such patterns are consistent with the magnet effect and may signal the overreaction of market participants, as described by Subrahmanyam (1994).

In summary, our exploratory analysis provides evidence for an increase in volatility and trading activity, as well as for a deterioration of liquidity around trading pauses. Solely based on these findings, however, we cannot claim that these effects are driven by trading pauses. They might as well be the result of a high-volatility market period with hectic trading activity. Hence, in order to make a causal claim, and avoid potential reverse causality, we need to control for effects that would have happened *without* trading pauses. This is carried out in the following section.

4 Do Trading Pauses Affect Volatility and Liquidity?

To address the first two research questions, we need to construct a set of control events from a time period when trading pauses are *not* in force, but market conditions are otherwise comparable, and the price movements of the same (or similar) stocks are sufficiently large, so that they would trigger a trading pause. Unfortunately, we cannot make use of cross-listings on other stock exchanges, as trading pauses have been implemented synchronously across all exchanges in the U.S. Likewise, investigating cross-listings in Europe is not feasible as trading hours have only a small overlap between the two continents. This would dramatically reduce the number of trading pauses in the sample and severely weaken the validity of our empirical results.

A viable option, however, is to rely on extreme price movements on NASDAQ from the pre-regulation period between January 2009 and May 2010, when trading pauses had not yet been implemented, but market conditions were largely similar. From this time period, we collect events when stock prices changed by more than 5% over five minutes, so that they would have triggered a trading pause, had trading pauses been in force at the time. The two-hour intervals around these events build a control group for each stock that triggered a trading pause in our sample. We match the control events by stock, while accounting for the direction of the underlying price movement. To avoid overlapping effects, we only consider control events that are at least 120 minutes apart from each other. Figure 7 shows the cumulative distribution of the five-minute price changes that are used as control group. The total number of extreme price movements is 1755, which results into nine matched events per trading pause on average.

To verify that market conditions in the pre-regulation (control) period are similar to those in the regulation (sample) period, we compare the evolution of the (log) bid-ask spread, the (log) trading volume and the (annualized) realized volatility⁶ between the two periods. Figure 21 in the Appendix illustrates that spreads, trading volumes and volatility behave quite similarly in both periods. This is true for both their cross-sectional averages and cross-sectional distributions. We formally test the equality of means and medians of the three variables between the two

⁶Realized volatility is estimated for each day in the entire time period (i.e. control and sample period), using one-hour mid-quote returns of paused stocks. Daily statistics are then expressed on an annual scale.

periods using paired t-tests and Wilcoxon signed-rank tests, and find no significant difference between the control and the sample periods (see Table 2).

Figure 7: Cumulative Distribution of Selected Pre-Regulation Price Changes

This figure shows the (cumulative) distribution of the 1755 five-minute stock price changes that we use as control group in our analysis. These control events originate from the time period between Jan 2009 and May 2010, and consist of price movements which are large enough (i.e. greater than 5%) in absolute value in order to trigger a trading pause, had this regulation been in force at the time. We separate between the distributions of downward (red) and upward (blue) price movements, which are matched with the direction of the underlying price movement and the respective stock.

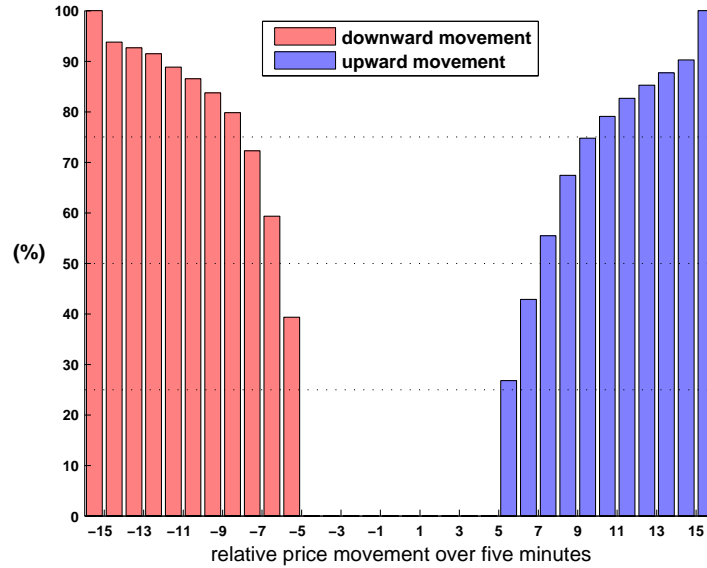


Table 2: Market Characteristics in the Control and Sample Periods

This table compares differences between the pre-regulation (control) and post-regulation (sample) period, from January 2009 to May 2010 and from June 2010 to June 2014, respectively. The comparison is based on three market measures: (log) bid-ask spread, (log) trading volume and (annualized) realized volatility. In the first two columns, we report the averages of these market measures in the pre-regulation and post-regulation periods. In the last two columns, the p-values of (parametric) paired t-tests and (non-parametric) Wilcoxon signed-rank tests are reported, comparing means and medians between the two periods, respectively.

	pre-regulation (control period)	post-regulation (sample period)	t-test (p-value)	Wilcoxon test (p-value)
(log) bid-ask spread	-4.044	-4.132	0.134	0.187
(log) trading volume	9.883	9.990	0.150	0.226
(annual) realized volatility	19.85%	23.44%	0.161	0.111

Using the constructed control group of extreme price movements, we perform a difference-in-differences analysis, comparing each trading pause to one (or more) matched events using a pooled regression. We run separate regressions for each liquidity and volatility variable analyzed in Section 3.2, generically denoted by ξ . This variable ξ is recorded at each time bin t , $t = 0, \dots, N$, of the two-hour interval around trading pause h , $h = 1, \dots, H$, and the matched control event k , $k = 1, \dots, |\mathcal{K}^h|$, with $|\mathcal{K}^h|$ denoting the number of control events matched with the h^{th} trading pause.⁷ Using this notation, the corresponding regressions are formulated as

$$\xi_{t,h} = \alpha_h + \sum_t (\beta_t + \beta_t^{\text{TP}}) D_t + \varepsilon_{t,h}, \quad (3)$$

$$\xi_{t,h,k} = \alpha_{h,k} + \sum_t \beta_t D_t + \varepsilon_{t,h,k}, \quad (4)$$

where $\varepsilon_{t,h}$ and $\varepsilon_{t,h,k}$ denote white noise error terms, coefficients α_h and $\alpha_{h,k}$ capture fixed effects, and D_t are dummy variables indicating the respective time bins. The distinct effect of trading pauses is captured by coefficient β_t^{TP} , which is identified by first taking differences *within* equations (3) and (4), yielding

$$\xi_{t,h} - \xi_{0,h} = \sum_t (\beta_t + \beta_t^{\text{TP}}) D_t + \varepsilon_{t,h}^*, \quad (3^*)$$

$$\xi_{t,h,k} - \xi_{0,h,k} = \sum_t \beta_t D_t + \varepsilon_{t,h,k}^*, \quad (4^*)$$

and then taking differences *between* equations (3^{*}) and (4^{*}), leading to the final regression equation

$$\tilde{\xi}_{t,h,k} = \sum_t \beta_t^{\text{TP}} D_t + \tilde{\varepsilon}_{t,h,k}, \quad (5)$$

where $\tilde{\xi}_{t,h,k} := (\xi_{t,h} - \xi_{0,h}) - (\xi_{t,h,k} - \xi_{0,h,k})$ and $\tilde{\varepsilon}_{t,h,k} := \varepsilon_{t,h}^* - \varepsilon_{t,h,k}^*$. Potential error autocorrelation is removed by applying the method of Prais and Winsten (1954). Furthermore, we allow for heteroskedasticity both across trading pauses and time by assuming an error covariance matrix

⁷The 0th time bin is defined as the bin one hour before, while the Nth time bin is defined as the bin one hour after the corresponding trading pause.

of the form

$$\Sigma = \begin{bmatrix} \mathbf{I}_{|\mathcal{K}^1|} \otimes \Sigma_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{|\mathcal{K}^2|} \otimes \Sigma_2 & \cdots & \mathbf{0} \\ \vdots & & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{I}_{|\mathcal{K}^H|} \otimes \Sigma_H \end{bmatrix},$$

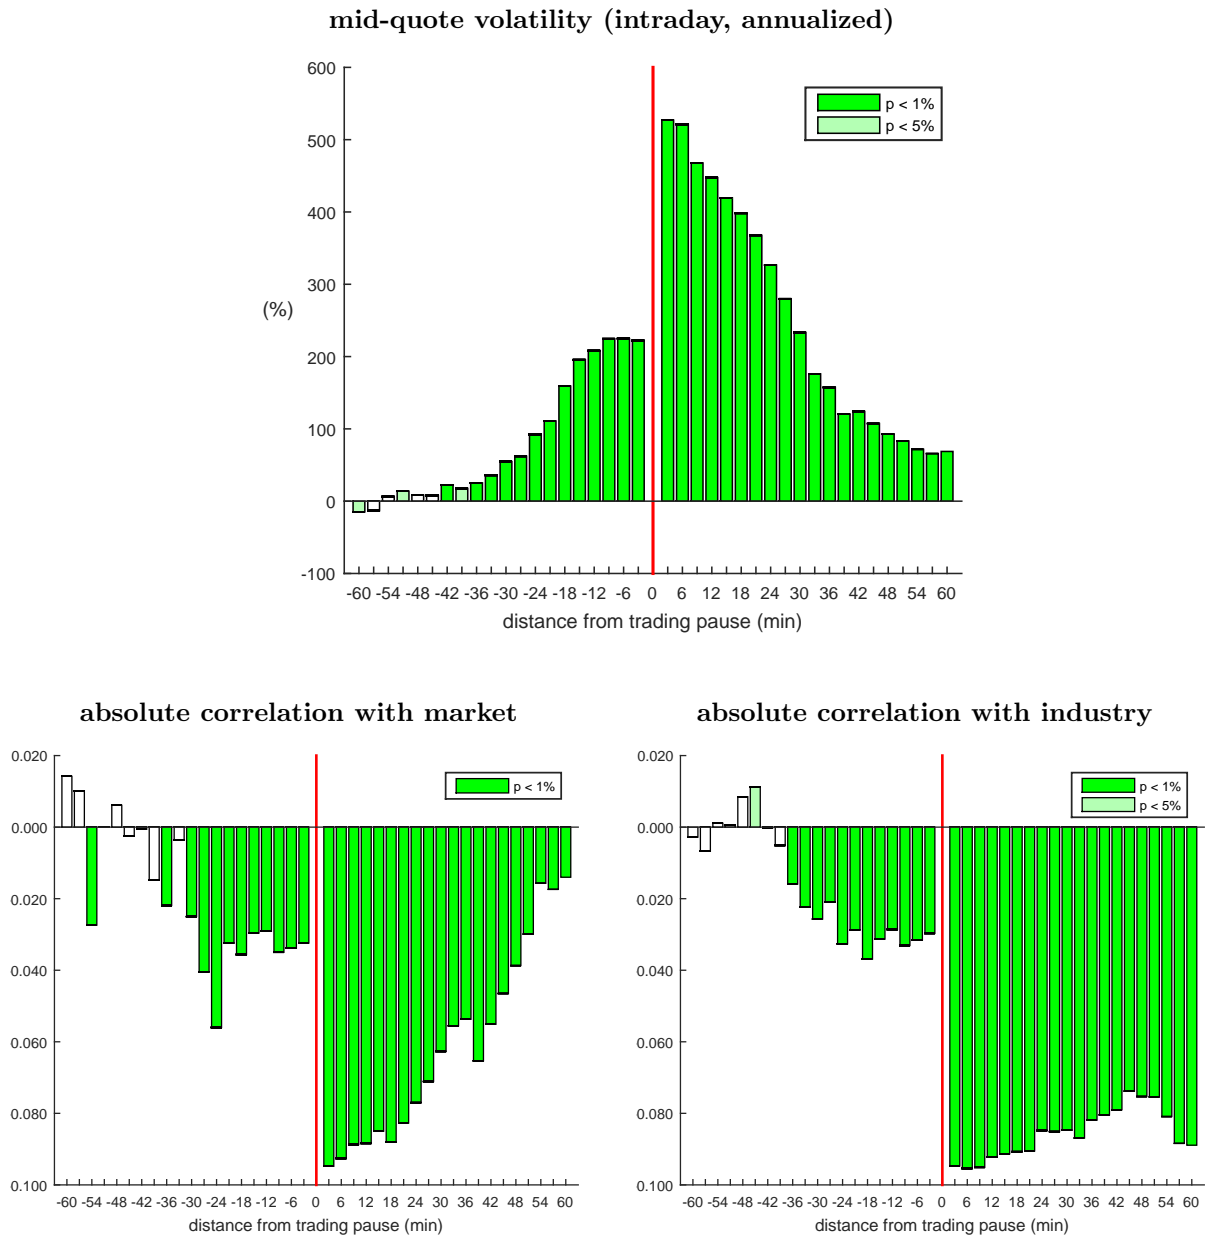
where \mathbf{I}_j denotes the identity matrix of dimension j , and the (pause-specific) error covariance matrix Σ_h is given by

$$\Sigma_h = \begin{bmatrix} \sigma_{h,1}^2 & 0 & 0 & 0 \\ 0 & \sigma_{h,2}^2 & 0 & 0 \\ 0 & 0 & \hat{\sigma}_{h,3}^2 & 0 \\ 0 & 0 & 0 & \sigma_{h,4}^2 \end{bmatrix} \otimes \mathbf{I}_{N/4}.$$

Running regression (5) for the intraday volatility and market/industry correlation measures, defined in Section 3.2, yields difference-in-differences estimates of β_t^{TP} for each 3-minute bin during the two-hour window around trading pauses. These coefficients represent the distinct (causal) effect of trading pauses on volatility and correlations, in comparison to the matched control events in the pre-regulation period. We graphically illustrate them in Figure 8. It turns out that trading pauses have a significant effect on intraday volatility. This effect is not only identifiable *after* trading pauses, but already *before*. In particular, during the 30-minute interval prior to the pause, volatility is significantly higher than in a situation where trading pauses cannot occur. Hence, the mere presence of such a regulation changes the behavior of volatility even before a trading pause is actually triggered. Hence, market participants seem to anticipate the possibility of a trading interruption, which is in line with the “magnet effect” hypothesis of Subrahmanyam (1994). After the trading interruption, we find that volatility is twice as high as in a market environment without trading pauses. This increased level persists for about 30 minutes. Therefore, consistently with the predictions of Fama (1989), trading pauses do not reduce volatility but merely postpone it.

Figure 8: Estimates of β_t^{TP} for Volatility and Correlation Measures

This figure depicts the estimates of β_t^{TP} , quantifying the effect of trading pauses on intraday volatility, (absolute) stock-to-market return correlation, and (absolute) stock-to-industry return correlation, estimated according to Section 3.2. Estimates of β_t^{TP} are calculated by means of difference-in-differences (DiD) analysis, based on the regression given in Equation (5), comparing 195 trading pauses to a control group of 1755 matched extreme price movements from the pre-regulation period (Jan 2009 to May 2010) over a series of 3-minute bins in the two-hour interval around trading pauses. The five-minute trading pause is marked by the vertical red line. The shading of the bars indicates the significance of the estimates at the 1% and 5% level, while transparent bars represent insignificant estimates.

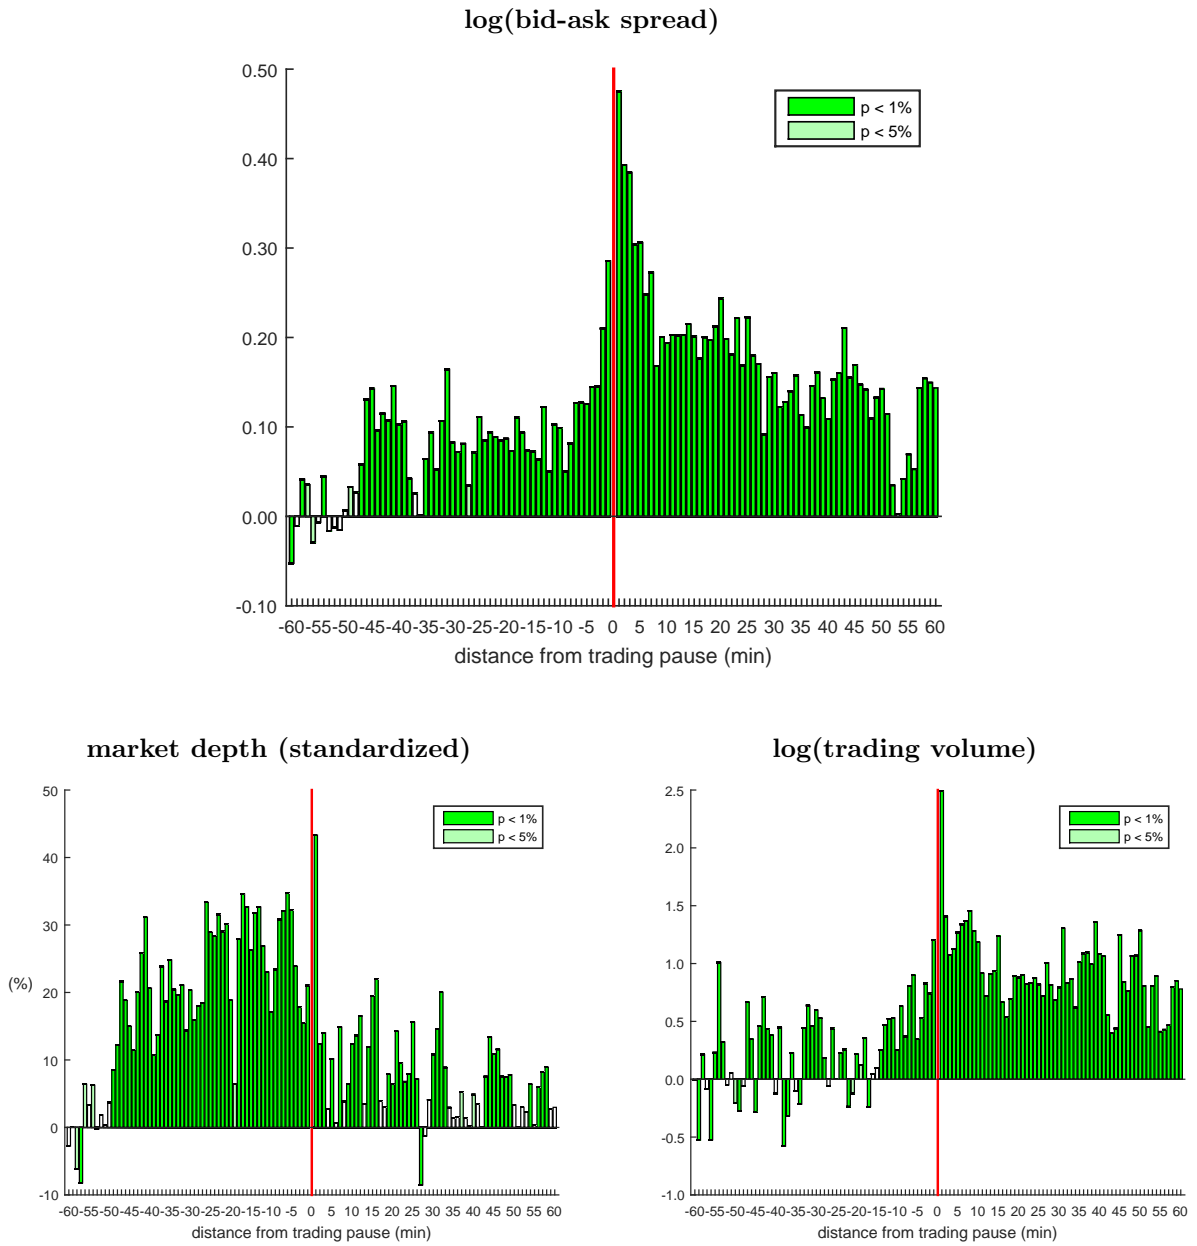


In addition, trading pauses have a significant negative effect on absolute high-frequency return correlations between the paused stock and the market, as well as between the stock and the respective industry. While we detect only a slightly negative effect on absolute correlations before the pause, this effect becomes significantly stronger thereafter. Hence, we find that trading pauses tend to push cross-correlations toward zero, weakening the connection between the paused stock and the rest of the market. One possible explanation for this phenomenon is that the occurrence of such a pause causes market participants to adjust their positions in the respective stock. Another reason for a higher idiosyncrasy could be the disruption of automated trading algorithms, which may react very differently to a trading pause than human traders, and handle paused stocks differently than the non-paused constituents of the stock market.

To analyze the effects of trading pauses on liquidity supply and demand, we perform difference-in-differences analyses for the 1-minute average bid-ask spread, average market depth and cumulative trading volume as defined in Section 3.2. The respective estimates of β_t^{TP} are shown in Figure 9. In line with the results above, we find significant effects already *before* trading pauses, showing that the “magnet effect” of a trading interruption is not only detectable for volatility but also for liquidity. In particular, we find that trading pauses significantly increase (log) bid-ask spreads. Approximately 45 minutes before the trading interruption, bid-ask spreads tend to be 10% higher than in a situation without trading pause regulation. When trading is resumed, this level is increased to approximately 20% and persists during the following hour. These patterns resemble the evolution of volatility during periods of highly unstable prices, and are presumably a manifestation of increased uncertainty on the market. The presence of a trading pause regulation makes liquidity providers obviously more hesitant to quote narrow spreads than they would normally do in volatile periods without trading pauses. Therefore, it seems that a trading pause mechanism is perceived as another source of uncertainty rather than as an additional layer of protection on the market. The fact that bid-ask spreads further widen after trading pauses and thus increase the costs of trading, indicates that trading pauses are not able to reduce uncertainty in the way postulated by Kyle (1988).

Figure 9: Estimates of β_t^{TP} for Market Liquidity Measures

This figure depicts the estimates of β_t^{TP} , quantifying the effect of trading pauses on average (log) bid-ask spreads, average market depth (standardized by its monthly average), and (log) cumulative trading volume, as defined in Section 3.2. Estimates of β_t^{TP} are calculated by means of difference-in-differences (DiD) analysis, based on the regression given in Equation (5), comparing 195 trading pauses to a control group of 1755 matched extreme price movements from the pre-regulation period (Jan 2009 to May 2010) over a series of 1-minute bins in the two-hour interval around trading pauses. The five-minute trading pause is marked by the vertical red line. The shading of the bars indicates the significance of the estimates at the 1% and 5% level, respectively, while transparent bars represent insignificant estimates.



Even though the presence of trading pauses makes liquidity providers to quote more cautiously, they do not seem to reduce their open positions in anticipation of a potential trading interruption. As shown in Figure 9, market depth on the first two order book levels significantly increases, which might be, however, partly driven by a widening of bid-ask spreads, causing a higher concentration of depth close to the best bid-ask quotes. We observe that market depth increases by approximately 25% prior to the pause. Hence, even though the presence of trading pause algorithms makes quoting activity less aggressive, it encourages liquidity suppliers to keep (or even increase) the size of their outstanding limit orders, regardless of an impending trading interruption. These findings are consistent with the argumentation of Kyle (1988), predicting that automatic trading interruptions protect investors and encourage liquidity suppliers to remain active even in turbulent times. In this sense, our results indicate that for liquidity suppliers, trading pauses may be perceived as safeguards. We observe that this effect remains after the trading pause, but becomes clearly weaker and partly insignificant. After the resumption of trading, liquidity supply is only slightly higher than in the control scenario without trading pauses. We therefore conclude that the post-pause effect on liquidity supply is rather weak.

The trading pause effect on liquidity *demand*, measured by the cumulative trading volume, is similar to the effect on volatility. We observe a gradual increase of trading activity approximately 20 minutes before the trading pause, which is probably a consequence of the “magnet effect” on volatility described earlier. We find that trading pauses make trading volume twice as large as during the matched control scenarios. Hence, the effect that trading pauses postpone and amplify volatility also manifests in increased trading activity. This is consistent with the well-known positive relationship between volatility and trading volume (Clark (1973); Epps and Epps (1976); Tauchen and Pitts (1983)).

Figure 10 provides insights into the effects of trading pauses on *directional* liquidity supply and demand, as reflected by the order book imbalance (IMBAL) and order flow (OFLOW) measures, calculated over 3-minute bins. Recall that the statistics are interacted with the sign of the underlying price change, such that positive (negative) values indicate order book and order flow movements in the same (opposite) direction. The estimates of β_t^{TP} for these two measures indicate that the presence of trading pauses has a significant effect on directional liquidity

supply and demand. The highly significantly negative coefficients after the pause indicate that both liquidity supply (IMBAL) and liquidity demand (OFLOW) are lower (or more negative) than in the control scenarios without trading pauses. Hence, when a large upward (downward) price movement results into a trading pause, traders tend to increase their selling (buying) activity. This is consistent with the “time-out” effect described by Brady (1988) and the “cool-off” effect postulated by Kyle (1988), arguing that a temporary trading suspension gives market participants time to process information and to revise outstanding positions.

In line with such a “time-out” effect, the directional shifts of liquidity supply according to Figure 10 might be caused by market participants who are willing to position themselves contrary to the price movement as long as they are protected by trading pauses. The latter limits their adverse selection risk and makes it more likely that counteracting forces become dominant – if the market has time to “cool off”. Note that these effects obviously exist under the presence of trading pause regulation, but not in the pre-regulation period. Hence, the mere existence of such a safeguard, protecting liquidity suppliers from adverse selection risk, makes market participants to behave differently in light of an imminent trading break. As indicated by the results above, they nevertheless refrain from quoting too aggressively, which results in wider spreads. As shown in a simple modification of the theoretical model by Glosten (1994) in Section 5.3, these shifts of liquidity supply might be (partly) due to informed market participants who counter-act against over-reactions of uninformed order flow, and thereby contribute to a (temporary) price stabilization after the trading pause. The analysis of quoting behavior during the pause in Section 5 provides additional support for this hypothesis.

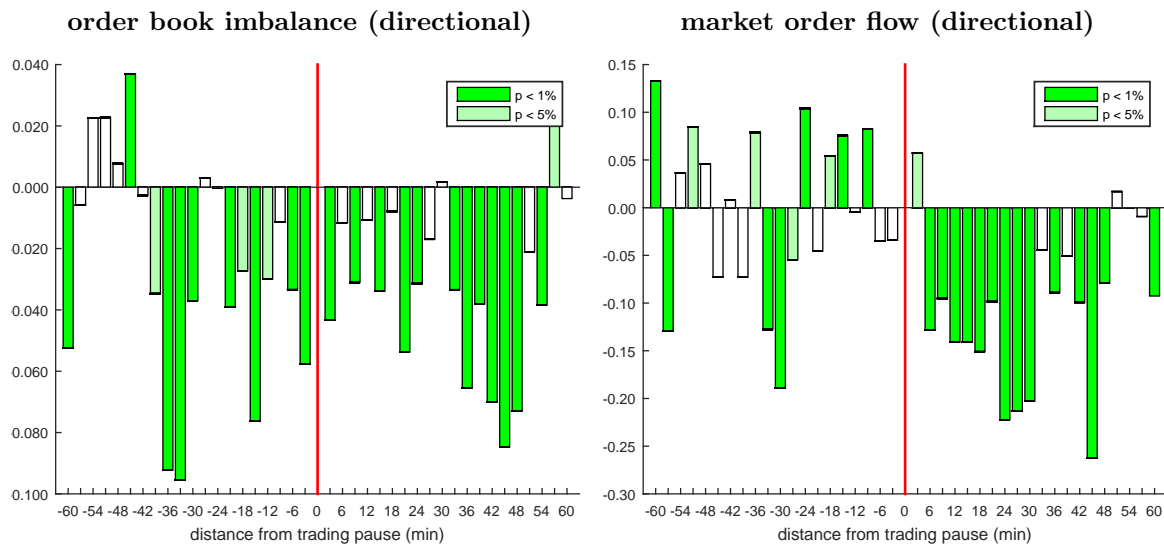
In summary, we find evidence for the causal effects of trading pauses on volatility and liquidity. Consistently with the “magnet effect” hypothesis, we observe that the evolution of volatility, bid-ask spreads and trading volume is significantly influenced by the anticipation of an imminent trading pause. The effects indicate that the presence of an automatic trading interruption mechanism does not calm market participants down, but rather acts as a source of additional uncertainty, inducing higher trading costs. Therefore, we do not find evidence for trading pauses acting as volatility circuit breakers and helping the market “cool off”. Instead, we observe that trading suspensions distort market dynamics, which materializes in the form of

higher volatility and transaction costs, as liquidity suppliers are hesitant to quote too close to the market.

On the other hand, our results also show that trading pauses create a protective layer through the anticipation of a “time-out” effect. The trading suspension allows market participants time to process information and revise their positions in hectic market periods. The mere existence of this safeguard makes liquidity suppliers changing their behavior when (the triggering of) a trading pause becomes likely. In such scenarios, (possibly better-informed) liquidity providers are incentivized to act as counterweights, and create positions in the opposite direction as the underlying price movement.

Figure 10: Estimates of β_t^{TP} for Market Asymmetry Measures

This figure depicts the estimates of β_t^{TP} , quantifying the effect of trading pauses on average (directional) order book imbalance and (directional) cumulative market order flow, as defined in Equations (1) and (2) in Section 3.2. Estimates of β_t^{TP} are calculated by means of difference-in-differences (DiD) analysis, based on the regression given in Equation (5), comparing 195 trading pauses to a control group of 1755 matched extreme price movements from the pre-regulation period (Jan 2009 to May 2010) over a series of 3-minute bins in the two-hour interval around trading pauses. The five-minute trading pause is marked by the vertical red line. The shading of the bars indicates the significance of the estimates at the 1% and 5% level, while transparent bars represent insignificant estimates.



5 Order Book Activity During Trading Pauses

Our data enables us to analyze order book activity and price formation, not only around but also *during* trading pauses. In this section, we address our third research question, and investigate to what extent market participants make use of the temporary trading suspension by revising their positions and contributing to price discovery.

5.1 Order Placement

Figure 20 in the Appendix shows changes in (standardized) limit order book depth around and during trading pauses.⁸ The left panel depicts changes in market depth during the last 5 minutes preceding and the first 5 minutes after the trading pause. For better visibility, the time period during the pause is shown separately in the right panel. The bar charts show the cumulative order book depth for 10 price bins, ranging from 90% to 110% around the current mid-quote.

We observe that order book depth declines rapidly right before the trading pause, and then gradually recovers during the five-minute suspension. Interestingly, this replenishment, however, is asymmetric, and depends on the direction of the underlying price movement triggering the pause. Accordingly, we observe selling pressure in the limit order book after upward price movements and buying pressure after downward movements. These patterns are in line with the findings in Section 4 on order book imbalances, and are consistent with the logic of Brady (1988) and Kyle (1988). Hence, a temporary suspension facilitates the recovery of prices from shocks by giving market participants time to revise positions, which foster a trend reversal. These results suggest that the limit orders placed during trading pauses are (to some extent) informative, in the sense that they contribute to a temporary price stabilization *after* the trading interruption.

Note that during the trading pause, outstanding limit orders cannot be matched. Consequently, liquidity providers do not face any adverse selection risk, as they cannot get picked up by better-informed (or faster) market participants. Betting on the “wrong” side of the market, however, can still be costly, as it entails a later re-positioning of limit orders. Such adjustments cause additional transaction costs as they induce a loss of priority in the order queue. Therefore, even

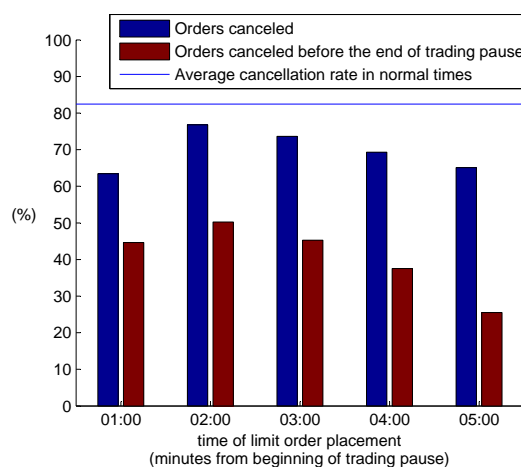
⁸To make effects comparable across stocks, we standardize limit order book depth by its one-month average, computed over the first twenty order book levels.

in the absence of adverse selection risk, these implicit costs enforce liquidity providers to place limit orders in a sustainable way, avoiding expensive re-positioning.

Indeed, this argumentation is consistent with the development of order cancellation ratios during the pause. Figure 11 shows the percentage of cancelled limit orders, which are deleted either during the pause or thereafter. Comparing the order cancellation ratios during trading pauses to cancellation ratios outside of trading pauses (depicted by the horizontal blue line in the figure), we conclude that the limit orders posted during trading pauses have longer holding times than those posted in normal times. While the average order cancellation ratio outside of trading pauses exceeds 80%, this ratio is around 70% for limit orders submitted during pauses. Moreover, we find that the proportion of limit orders which are canceled during the pause (i.e. before the resumption of trading) is even lower, ranging between 30% and 50%.⁹ Hence, we can conclude that more than 50% of all limit orders posted during a trading pause “survive” the interruption, and thus are indicative for post-pause price levels. Therefore, we argue that intra-pause limit orders reveal trading intentions, and are not (merely) used to “test” the market.

Figure 11: Cancellation Ratios of Intra-Pause Limit Orders

This figure shows the percentage of limit order volume that is posted during trading pauses and canceled later, either during the pause (red bars) or after the pause (blue bars). The percentages are computed for orders arriving during any of the 195 five-minute trading pauses in our sample. For the sake of comparison, the horizontal blue line represents the average cancellation ratio of limit orders posted in normal times (i.e. trading days without trading pauses).



⁹The gradual decline of this ratio is arguably attributable to the simple effect that as we get closer to the end of the trading pause, there is less time for traders to change their strategies and cancel orders.

5.2 Price Convergence

Since order execution is suspended during the trading pause, the market is not cleared, which potentially results in overlapping bid and ask quotes. Therefore, in order to analyze price convergence during trading pauses, we (artificially) clear the market by matching executable limit orders against each other, and compute the (hypothetical) mid-quote price at each instant.

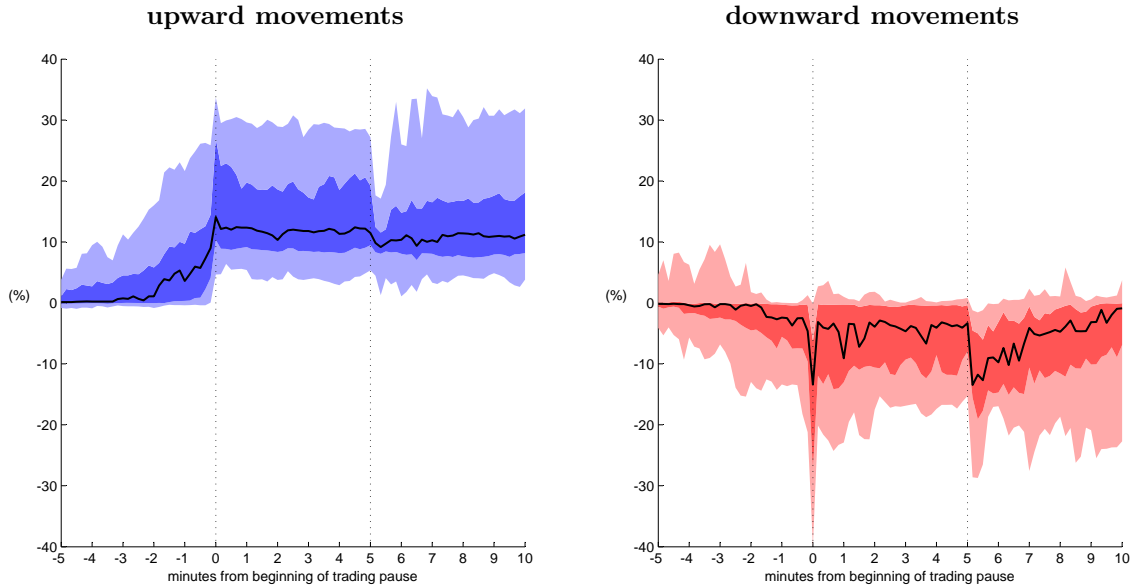
Figure 12 shows the resulting cross-sectional variation of the cumulative intra-pause mid-quote returns. Differentiating between trading pauses triggered by upward and downward price movements, we observe that “upward” and “downward” pauses exhibit different price trajectories. In the case of upward moving prices, prices slightly reverse at the beginning of the trading pause, but quickly stabilize and maintain their level until the end of the break. Then, in the first minute after the resumption of trading, there is a slight downward correction towards a stable price level, which, however, is significantly higher than before the trading pause. The observed price evolution is therefore consistent with the revelation of positive news, as also documented in the empirical work of Bhattacharya and Spiegel (1998).

In contrast, the instantaneous corrective forces of trading pauses after downward price movements seem to be weaker without revealing a clear price pattern. Even though trading pauses alleviate the downward price pressure and mid-quote prices stabilize during the suspension, they tend to revert back to their initial level when the pause expires. After this drop, we see a gradual convergence to a stable price, which is around the same level as prior to the trading pause. Consistently with the argumentation of (Harris, 1998) and corresponding discussions in Chapter 5.3, this pattern suggests a rather transitory shock of the stock price.

Two aspects, however, should be noted. Firstly, as shown in Figure 12, the cross-sectional variation, and thus the (empirical) confidence intervals, of the mid-quote price paths are relatively large. Interpretation of mean effects should therefore be done with caution. Secondly, our inference rests on the assumption that the hypothetical (intra-pause) mid-quote prices are representative for the “true” price level during the trading pause. We presume that market participants have a notion of the implied mid-quote level and keep track of it during a trading pause.

Figure 12: Cumulative Stock Returns Around and Within Trading Pauses

This figure shows the cross-sectional distribution of cumulative mid-quote returns before, during and after the 195 trading pauses in our sample. The cumulative returns are calculated every 15 seconds, and then scaled to a daily level. Within each pause, mid-quote returns are constructed by artificially clearing the market. The shaded bands in the charts correspond to the 16th, 33rd, 50th (black line), 67th and 84th percentiles of the empirical cross-sectional distribution based on 195 trading pauses in the sample. The two vertical dotted lines represent the beginning and the end of the five-minute pause.



To gain a deeper understanding of the price convergence during trading pauses, we perform a similar analysis as in Corwin and Lipson (2000). By means of regression analysis, we quantify to what extent a temporary suspension facilitates convergence to a “stable” price level. If trading pauses alleviate market tensions and accelerate information processing, then a large part of price stabilization should take place during the pause itself. Such an analysis can shed light on the quality of price discovery during trading pauses, but it crucially depends on the identification of the “stable” price. In the given context, we consider the first post-pause price level which persists without a statistically significant trend for at least ten minutes as (temporarily) “stable”. After each trading pause, we determine the earliest time point when this condition is satisfied, and denote it by $T(h)$. If price stabilization cannot be detected in the one-hour window after the trading pause, we deem the price level “non-convergent”, and exclude it from the following regression analysis. This leaves us with a sample of 158 trading pauses out of the 195 trading pauses used in Section 4.

By computing log returns from the beginning of the h^{th} trading pause to $T(h)$ and regressing them on intra-pause returns of different lengths, we quantify how much of the price change toward the “stable” price level is already realized during the trading pause. This leads to the regression equation

$$R_{[\text{H0}(h), \text{H1}(h)+T(h)]} = \beta_t R_{[\text{H0}(h), \text{H0}(h)+t]} + \varepsilon_{t,h}, \quad h = 1, \dots, 195, \quad (6)$$

where $R_{[a,b]}$ denotes the log return from a to b , h indexes the trading pauses, $H0$ and $H1$ represent the beginning and end of the h^{th} pause, respectively, and $\varepsilon_{t,h}$ denotes a white noise error term. The index t refers to a given time point during the trading pause. We run different regressions by varying t from 0.5 to 5.0 minutes yielding corresponding estimates of β_t .

Since the unconditional expectation of short-term log returns is very close to zero, it is easily justifiable to omit the intercept term from the regression equation. Then, β_t naturally measures the portion of the “stable” return realized until t through the trading pause. The benchmark case $\beta_t = 1$ implies perfect price discovery, in the sense that the “stable” price level at $T(h)$ is reached already at t . Accordingly, if $\beta_t \neq 1$, the intra-pause price level at time t systematically undercuts ($\beta_t > 1$) or exceeds ($\beta_t < 1$) the “stable” price level.

The estimation of β_t , however, is hindered by the overlap of $R_{[\text{H0}(h), \text{H1}(h)+T(h)]}$ and $R_{[\text{H0}(h), \text{H0}(h)+t]}$, causing simultaneity and rendering OLS estimates inconsistent. We address this problem by splitting $R_{[\text{H0}(h), \text{H1}(h)+T(h)]}$ into two parts

$$R_{[\text{H0}(h), \text{H1}(h)+T(h)]} = R_{[\text{H0}(h), \text{H0}(h)+t]} + R_{[\text{H0}(h)+t, \text{H1}(h)+T(h)]},$$

and rearranging the initial regression equation to

$$R_{[\text{H0}(h)+t, \text{H1}(h)+T(h)]} = (\beta_t - 1)R_{[\text{H0}(h), \text{H0}(h)+t]} + \varepsilon_{t,h}.$$

Even though this proceeding tackles the problem of overlapping returns, it still does not allow us to identify the causal effect of trading pauses. Hence, similarly to the procedure in Section 4, we isolate the trading pause effect from the effect of the underlying price movement

by using not only the trading pauses themselves but also the corresponding control group of extreme price movements in the pre-regulation period. This leads to a regression utilizing the *pooled* sample of control group events and pause events. With i indexing both trading pauses and control events, the final regression specification is

$$R_{[\text{H0}(i)+t, \text{H1}(i)+T(i)]} = \left[(\beta_t - 1) + \beta_{t,U}U_i + (\beta_{t,P} + \beta_{t,PU}U_i)P_i \right] R_{[\text{H0}(i), \text{H0}(i)+t]} + \varepsilon_{t,i}, \quad (7)$$

where P_i and U_i are dummy variables taking the value one if the i^{th} event is a trading pause and is associated with an upward price movement. Therefore, β_t and $(\beta_t + \beta_{t,U})$ measure the strength of price discovery t minutes after the occurrence of a downward or upward price movement in the control group, respectively. Likewise, price discovery during trading pauses triggered by downward or upward price movements is measured by the sums $(\beta + \beta_{\text{H}})$ and $(\beta + \beta_{\text{U}} + \beta_{\text{H}} + \beta_{\text{HU}})$, respectively.

The regression results for different values of t are summarized in Table 3. We test the (sums of) coefficient estimates against the benchmark $\beta = 1$, indicating perfect price discovery. As shown by the β_t estimates, which do not significantly differ from 1 for any t , price discovery is nearly perfect after extreme downward price movements in the control group. Hence, the price levels reached shortly after downward price shocks are representative of the “stable” price level. Therefore, in these situations, price stabilization is naturally implied by market forces, without the intervention of trading pauses.

In contrast, extreme upward price movements without trading interruptions are typically followed by further price increases over the next four minutes. This is identifiable from $(\beta + \beta_{\text{U}})$, which significantly exceeds 1 for $t < 4$, implying that prices systematically undershoot the “stable” price level. However, if a trading pause is triggered, price convergence and information processing is accelerated, and pricing errors are only detectable in the first 30 seconds after the beginning of the pause. This last observation follows from $(\beta_t + \beta_{t,U} + \beta_{t,P} + \beta_{t,PU})$ being significantly different from 1 only for $t = 0.5$.

Table 3: Price Discovery Regressions


This table shows the parameter estimates from regression (7), fitted for different time points t ranging over a 30sec grid. Estimation is performed using the sub-sample of 158 trading pauses and their (matched) control events that reach a stable price level in the hour after the extreme price movement. Together with the control events, this gives a total sample size of 1889. The estimated (sums of) coefficients are tested against one using standard errors robust to pause-level heteroskedasticity. The significance of t-statistics is denoted by *** : $p < 0.01$, ** : $p < 0.05$ and * : $p < 0.10$.

t (minutes)	control events		trading pauses		R ²
	β_t	$\beta_t + \beta_{t,U}$	$\beta_t + \beta_{t,P}$	$\beta_t + \beta_{t,U} + \beta_{t,P} + \beta_{t,PU}$	
0.5	1.132	1.172***	0.994	1.736***	0.83
1.0	1.046	1.088**	0.973	1.000	0.88
1.5	1.006	1.068**	0.948	1.001	0.89
2.0	0.997	1.089***	0.944	1.001	0.90
2.5	1.028	1.071***	0.969	0.999	0.90
3.0	1.025	1.065***	0.983	0.998	0.91
3.5	0.991	1.055**	0.998	0.995	0.91
4.0	1.003	1.036	0.999	0.996	0.91
4.5	1.010	1.024	1.022	1.016	0.92
5.0	1.023	0.995	1.009	1.002	0.92

This analysis provides some evidence for trading pauses contributing to temporary price stabilization. It should be noted, however, that the analysis discards 37 trading pauses for which we cannot find a “stable” price in the one hour following the pause. Our findings nonetheless indicate that trading pauses facilitate price recovery after temporary shocks, and contribute to temporary price stabilization. The results are consistent with the argumentation of Harris (1998), claiming that if price uncertainty is driven by noise, a trading interruption may reduce information asymmetry and enhance price discovery by encouraging informed traders to enter the market. These results are further confirmed by Figure 13, which shows the evolution of high-frequency mid-quote volatility during trading pauses. We observe that intra-pause (implied) mid-quotes quickly stabilize after the beginning of the pause, resulting in a volatility level which is significantly below the level prior to, and particularly after, the pause.

Figure 13: Median Annualized Realized Volatility During Trading Pauses

This figure shows the median annualized realized volatility of the mid-quote before, during and after trading pauses. The median is calculated over the cross-section of 195 trading pauses in our sample. During trading pauses, we synthetically clear the market by matching limit orders, since the limit order book may contain overlapping buy and sell limit orders, as order execution is suspended. Realized variance is calculated over a time grid of disjoint 15 second intervals, by taking the sum of squared (log) returns within the respective interval. The vertical dotted lines (at 0 and 300 seconds) represent the beginning and end of the five-minute trading pause, respectively.



5.3 Information Asymmetry in the Limit Order Book

To better understand how differences in traders' information affect bid-ask spreads and limit order submissions, we present a stylized limit order book model based on Glosten (1994). The purpose of this section is to demonstrate how the asymmetric limit order book patterns observed in Figure 20 may result from differently informed liquidity supplier groups, and why *informed* liquidity suppliers are likely to position themselves against the direction of the underlying price movement.

As a starting point, we assume the presence of uninformed *liquidity takers* submitting market orders and prioritizing execution immediacy over trading profit. The flow of these market orders is modeled as a normally distributed random variable $q \sim N(\mu, \sigma^2)$, which is centered around the (aggregated) price expectation of liquidity takers. Moreover, there are *liquidity providers*, who have expectations about the incoming market order flow q , and place their limit orders

accordingly, as long as their expected profits are non-negative. By balancing the risk of non-execution against the risk of (adversely) being picked off, liquidity providers expect a profit Π on the marginal limit order placed at price level p in the limit order book

$$\Pi(p) = \mathbb{E}[\mathbb{1}(Y(p)) \cdot |p - v|] - c, \quad (8)$$

where $Y(p)$ is the cumulative order book depth at price level p , v is the true (fundamental) value of the stock, c denotes the (unit) cost of order placement, and $\mathbb{1}$ indicates the execution of the limit order.¹⁰ Expressing equation (8) for the bid side of the limit order book, we obtain

$$\Pi(b) = P(Y(b)) \cdot [\mathbb{E}(v|Y(b) \leq q) - b] - c,$$

where $P(Y(b))$ is the probability for the cumulative order book depth at bid price level b getting executed by incoming market orders. By assuming the random market order flow q to be independent of the fundamental value v , the above expression simplifies to

$$\Pi(b) = \Phi_q(Y(b)) \cdot [\mathbb{E}(v) - b] - c, \quad (9: \text{BID})$$

for the bid side. Analogously, for the ask side, we have

$$\Pi(a) = [1 - \Phi_q(Y(a))] \cdot [a - \mathbb{E}(v)] - c, \quad (9: \text{ASK})$$

with a denoting the ask price level, and Φ_q representing the cumulative density function of the incoming market order flow q .

Setting the marginal profit equations (9: BID) and (9: ASK) equal to zero, they can be solved for the (cumulative) limit order book depth Y at any price level. This yields the equilibrium limit order book depth $Y^*(p)$ in dependence of the expectation about the fundamental value v , the distribution of market orders, and the cost of order placement. For the bid side, we therefore

¹⁰Whether an outstanding limit order gets executed, ultimately depends on the incoming market order flow, generated by uninformed liquidity takers.

obtain $Y^*(p)$ as

$$Y^*(p) = \begin{cases} \Phi_q^{-1} \left(\frac{c}{E(v)-b} \right) & \text{for } b < E(v) - \frac{c}{\Phi_q(0)}, \\ 0 & \text{otherwise,} \end{cases} \quad (10: \text{ BID})$$

and for the ask side

$$Y^*(p) = \begin{cases} \Phi_q^{-1} \left(1 - \frac{c}{a-E(v)} \right) & \text{for } a \geq E(v) + \frac{c}{1-\Phi_q(0)}, \\ 0 & \text{otherwise.} \end{cases} \quad (10: \text{ ASK})$$

Note that if the price of the buy (sell) limit order is too high (low), the risk that the limit order will be unprofitable (due to unfavorable price changes) is too high for the liquidity maker to break even. This creates a bid-ask spread in the order book, where no limit order is entered.

This setting corresponds to a simplified version of Glosten (1994), which yields a bid-ask spread, based on the differences between liquidity providers and liquidity takers. Now, we slightly modify this model by making a distinction between informed and uninformed liquidity providers, who have different beliefs about the fundamental value v of the stock: (i) Informed liquidity providers (e.g. professional investors) are assumed to have an unbiased forecast $E_i(v)$ of the fundamental value v . (ii) Uninformed liquidity providers (e.g. market makers), have a prediction $E_u(v)$ of v , which may be potentially biased¹¹. Moreover, we assume that uninformed liquidity providers dominate informed liquidity providers in terms of trading technology in the sense that they face lower transaction costs (i.e. $c_u < c_i$).

In this setting, equations (10: BID) and (10: ASK) can still be used to derive the equilibrium order book depth quantities $Y_i^*(p)$ and $Y_u^*(p)$ for both types of liquidity providers, given their (potentially different) beliefs about v and trading costs c . Taking the maximum of these quantities yields the aggregate limit order book depth at any price level p given by


$$Y^*(p) = \max\{Y_i^*(p), Y_u^*(p)\}. \quad (11)$$

¹¹For the sake of simplicity, let uninformed liquidity providers and liquidity takers have homogeneous expectations, i.e. $E_u(v) \equiv \mu$.

By refraining from endogenizing liquidity demand and ruling out strategic interactions between the two groups of liquidity providers, the model is undoubtedly very simplistic. While further extensions are beyond the scope of this paper, it nevertheless fulfills the purpose of providing some underpinning of the empirical evidence provided in the previous subsection. In particular, employing (10: BID), (10: ASK) and (11), we provide numerical examples of the limit order book in three different market scenarios. In the baseline scenario, we assume that the expectation of informed and uninformed traders are identical. This situation implies a symmetric limit order book, as visualized in Figure 14 for $E_i(v) = E_u(v) = \mu = 100$. As uninformed liquidity providers (i.e. market makers) have a comparative cost advantage, they can afford posting more limit orders than informed liquidity providers at any given price level, and thus ultimately determine the structure of the limit order book.

Figure 14: Simulated Limit Order Book Without Information Asymmetry

This figure shows the limit order book implied by our model in periods without information asymmetry. The left (right) side depicts the cumulative buy (sell) limit order quantity. The bars represent limit orders posted by uninformed liquidity makers. In the given setting, there are no limit orders posted by informed liquidity makers in excess of those represented by the blue bars.

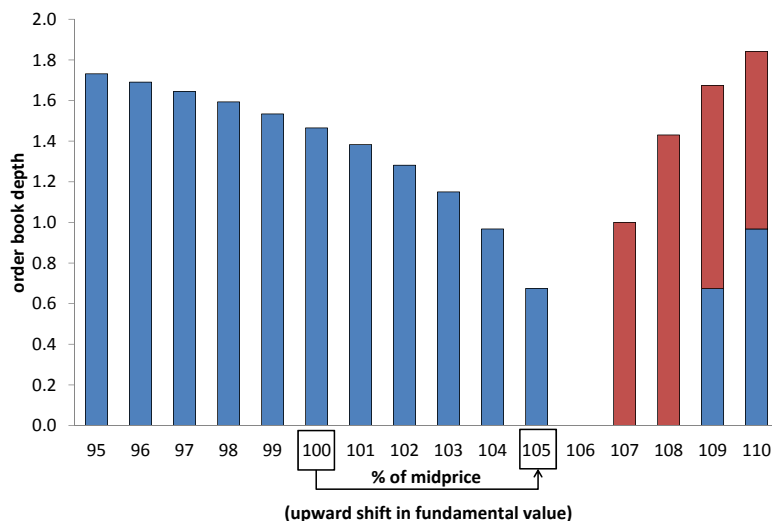


In the second scenario, we imitate the price pattern observed around upward trading pauses (see Figure 12), and assume a positive shock (e.g. “good news”) to the market, which permanently increases the stock’s fundamental value v by 5%. The shock is reflected in the expectations of

both informed and uninformed liquidity suppliers. The latter, however, are assumed to overreact. Accordingly, we let $v = E_i(v)$ to change from 100 to 105, but assume that $E_u(v) = \mu$ changes from 100 to 107. According to Figure 15, this scenario implies an asymmetric limit order book. Expecting that the fundamental value v equals 105, but uninformed market order flow is centered around 107, informed liquidity providers (i.e. professional investors) make use of their informational advantage and place aggressive limit orders in excess of uninformed limit orders on the sell side of the market. This implies a (corrective) downward pressure on the mid-quote price as well as a narrowing of the bid-ask spread compared to Figure 14.

Figure 15: Limit Order Book After Permanent Positive Shock

This figure shows the limit order book implied by our model after a 5% permanent increase of the fundamental value. The left (right) side depicts the cumulative buy (sell) limit order quantity. Blue bars represent limit orders posted by uninformed liquidity providers, while red bars represent limit orders posted by informed liquidity providers, *in excess of* the blue bars. In this numerical example, uninformed traders overreact to the upward shock, and position themselves around 107, instead of 105. In response to this, informed liquidity providers place ask limit orders in excess of uninformed limit orders on the sell side of the market, which narrows the bid-ask spread and causes a (corrective) downward pressure on the mid-quote price.




In the third scenario, we imitate the price pattern observed around downward trading pauses (see Figure 12), and model a negative shock (e.g. “bad news”) to the market. In contrast to the scenario above, we assume, however, that this shock is temporary and has no effect on the stock’s fundamental value v . Accordingly, the shock does not change the expectations of informed traders,

while uninformed traders are assumed to overreact. Hence, we simulate the model with v remaining constant at $v = E_i(v) = 100$, while $E_u(v) = \mu$ changes from 100 to 98. As illustrated by Figure 16, also this scenario implies an asymmetric limit order book. Now, uninformed liquidity is centered around 98, while informed liquidity providers place more aggressive limit orders on the buy side of the market. This yields upward price pressure and again a tightening of the bid-ask spread.

Figure 16: Limit Order Book After Transitory Downward Shock

This figure shows the limit order book implied by our model after a temporary downward shock to the fundamental stock value. The left (right) side depicts the cumulative buy (sell) limit order quantity. Blue bars represent limit orders posted by uninformed liquidity providers, while red bars represent limit orders posted by informed liquidity providers, *in excess of* the blue bars. In this example, the downward shock does not change the fundamental value. However, uninformed traders overreact to the shock and position themselves around 98, instead of 100. In response to this, informed liquidity providers place bid limit orders in excess of uninformed limit orders on the sell side of the market, which narrows the bid-ask spread and causes a (corrective) upward pressure on the mid-quote price.



Hence, irrespective of the nature of the underlying shock, we observe that in both previous scenarios, the resulting limit order book becomes asymmetric. In fact, the greater the difference in traders' expectations about the stock's fundamental value, the more pronounced this effect becomes. Hence, Glosten's model predicts that the order book asymmetry observed in Figure 20 results from an overreaction of uninformed traders, which forces informed liquidity suppliers to position themselves against the direction of the underlying price movement. Although we cannot empirically identify whether these counter-acting forces indeed originate from superior

information, our results show that they exist, and trading pauses help them to become effective and to contribute to price stabilization.

Therefore, by suspending order execution, trading pauses create a market environment where (informed) market participants are protected from being picked-off by more cost-efficient traders in an over-shooting market. Under such a protection, these liquidity suppliers have a stronger incentive to reveal their superior knowledge and post “price-correcting” limit orders on the opposite side of the market. These findings support the results in Section 3.2 on the effects of trading pauses on directional shifts in liquidity supply.

6 Conclusions

Trading pauses are important regulatory measures to maintain price stability and establish safeguards in case of operational problems, which may arise from automated trade execution, or the malfunctioning of high-frequency trading algorithms. The effect of these safeguards on market quality, however, is controversially discussed. In this paper, we provide empirical evidence on the effects of trading pauses introduced in the U.S. after the Flash Crash on May 2010. Utilizing high-frequency limit order book data from NASDAQ, we analyze these effects on intraday volatility, liquidity, as well as price discovery. The distinct effects of trading pauses are separated from the effects of the underlying price movements by matching trading pauses to a control group of comparable extreme price movements from the pre-regulation period. Moreover, we provide insights into the quoting behavior of market participants during trading pauses, and shed some light on whether and how price stabilization during the interruption is achieved.

Based on our findings we can draw two major conclusions. First, the presence of an automatic trading suspension mechanism makes market participants to react differently in periods of extreme price movements. In these cases, a trading pause serves as a safeguard, limiting the adverse selection risk of liquidity providers who post limit orders against the current price trend. Trading pauses thus create an additional layer of protection, which makes it easier for liquidity supply to shift to the opposite side of the market and slow down (or reverse) price trends. Our results show that it is the mere presence of trading pauses, rather than their actual length, which makes

them effective. Second, trading interruptions obviously distort price dynamics, which manifests in higher volatility and wider bid-ask spreads. In line with the “magnet-effect”, postulated by Subrahmanyam (1994), we observe that both volatility and bid-ask spreads increase already before trading pauses are triggered and stay elevated for at least one hour thereafter. Hence, trading pauses do *not* act as volatility circuit breakers, but rather as volatility amplifiers.

Therefore, regulators face a trade-off between the benefits of trading pauses, in terms of their function as a safeguard, and their downsides, in terms of adverse effects on volatility and transaction costs. Balancing this trade-off could require re-considering the trading pause algorithm and to link it more directly to prevailing market conditions. On the one hand, local volatility should be explicitly taken into account, ensuring that trading pauses are triggered by local price trends, rather than by overall volatility. On the other hand, our findings suggest that the length of trading pauses could be shortened without overly restricting their function. This way, the accumulation of uncertainty during trading pauses could be kept to a minimum, which may further help to reduce the post-pause volatility effect.

References

- David Abad and Roberto Pascual. Switching to a temporary call auction in times of high uncertainty. *Journal of Financial Research*, 33(1):45–75, 2010.
- Uttpal Bhattacharya and Matthew Spiegel. Anatomy of a market failure: NYSE trading suspensions (1974–1988). *Journal of Business and Economic Statistics*, 16(2):216–226, 1998.
- Markus Bibinger, Nikolaus Hautsch, Peter Malec, and Markus Reiss. Estimating the quadratic covariation matrix from noisy observations: local method of moments and efficiency. *Annals of Statistics*, 42(4):80–114, 2014.
- Nicholas F Brady. Report of the presidential task force on market mechanisms, 1988. U.S. Presidential Task Force on Market Mechanisms.
- Jonathan Brogaard and Kevin Roshak. Prices and price limits. *Working paper*, 2015.
- U.S. CFTC and U.S. SEC. Findings regarding the market events of May 6, 2010, September 2010. URL <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>. Report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues.
- David D Cho, Jeffrey Russell, George C Tiao, and Ruey Tsay. The magnet effect of price limits: evidence from high-frequency data on Taiwan Stock Exchange. *Journal of Empirical Finance*, 10(1):133–168, 2003.
- William G Christie, Shane A Corwin, and Jeffrey H Harris. NADSDAQ trading halts: the impact of market mechanisms on prices, trading activity, and execution costs. *The Journal of Finance*, 57(3):1443–1478, 2002.
- Peter K Clark. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, pages 135–155, 1973.
- Shane A Corwin and Marc L Lipson. Order flow and liquidity around NYSE trading halts. *The Journal of Finance*, 55(4):1771–1805, 2000.

- Thomas W Epps and Mary Lee Epps. The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. *Econometrica*, pages 305–321, 1976.
- E F Fama. Perspectives on October 1987, or, What did we learn from the crash? In R W Kamphuis, editor, *Black Monday and the future of financial markets*, pages 71–82. Mid American Institute for Public Policy Research, Inc. and Dow Jones-Irwin, Inc., 1989.
- Lawrence R Glosten. Is the electronic open limit order book inevitable? *The Journal of Finance*, 49(4):1127–1161, 1994.
- Bruce C Greenwald and Jeremy C Stein. Transactional risk, market crashes, and the role of circuit breakers. *Journal of Business*, pages 443–462, 1991.
- Lawrence Harris. Circuit breakers and program trading limits: What have we learned? *Brookings-Wharton Papers on Financial Services*, 63, 1998.
- Albert S Kyle. Trading halts and price limits. *Review of Futures Markets*, 7(3):426–434, 1988.
- Charles Lee, Mark J Ready, and Paul J Seguin. Volume, volatility, and New York Stock Exchange trading halts. *The Journal of Finance*, 49(1):183–214, 1994.
- Ananth Madhavan. Trading mechanisms in securities markets. *The Journal of Finance*, 47(2):607–641, 1992.
- Sigbert J Prais and Christopher B Winsten. Trend estimators and serial correlation. Technical report, Cowles Commission discussion paper, 1954.
- Markus Reiss. Asymptotic equivalence for inference on the volatility from noisy observations. *Annals of Statistics*, 39(2):772–804, 2011.
- U.S. SEC. Proposed rules for clearly erroneous trades, June 2010. URL <http://www.sec.gov/news/press/2010/2010-104.htm>. Press Release No. 2010-104.
- U.S. SEC. National market system plan to address extraordinary market volatility, May 2011. URL <http://www.sec.gov/rules/sro/nms/2011/34-64547.pdf>. Release No. 34-64547.

U.S. SEC. SEC approves proposals to address extraordinary market volatility in individual stocks and broader stock market, June 2012. URL <http://sec.gov/news/pressrelease/2012-2012-107.html>. Press Release No. 2012-107.

U.S. SEC. SEC issues order modifying and extending the pilot period for the national market system plan to address extraordinary market volatility, April 2016. URL <http://sec.gov/news/pressrelease/2016-75.html>. Press Release No. 2016-75.

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

George E Tauchen and Mark Pitts. The price variability-volume relationship on speculative markets. *Econometrica*, pages 485–505, 1983.

Kai Zimmermann. Price discovery in European volatility interruptions. *Working paper*, 2013.

Appendix

Figure 17: Volatility and Correlation Measure Quantiles Around Trading Pauses

This figure shows how the cross-sectional quantiles of intraday mid-quote volatility, (absolute) stock-to-market correlation, and (absolute) stock-to-industry correlation paths change in the two-hour interval around trading pauses. The plotted variables are defined in Section 3.2. The shaded bands in the chart correspond to the 16th, 33rd, 50th (black line), 67th and 84th percentiles, calculated for a series of 3-minute bins over the cross-section of 195 trading pauses in our sample. The left (right) part of the chart represents the pre-pause (post-pause) period, while the five-minute trading pause is denoted by the vertical red line.

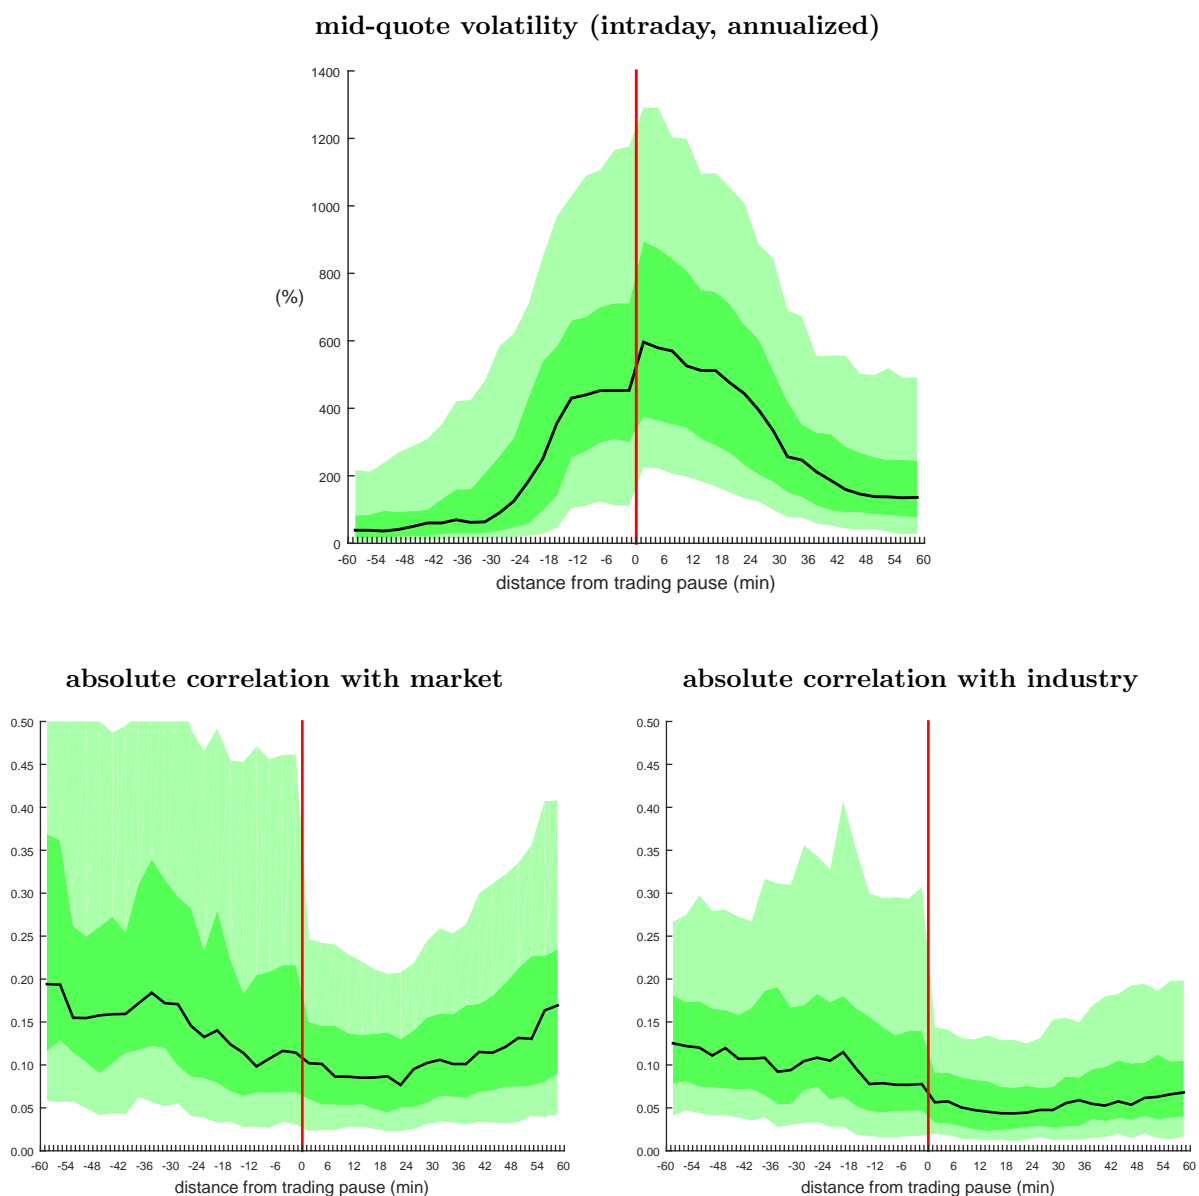


Figure 18: Market Liquidity Measure Quantiles Around Trading Pauses

This figure shows how the cross-sectional quantiles of average (log) bid-ask spreads, average two-level market depth (standardized by the monthly average), and (log) cumulative trading volume paths change in the two-hour interval around trading pauses. The plotted variables are defined in Section 3.2. The shaded bands in the chart correspond to the 16th, 33rd, 50th (black line), 67th and 84th percentiles, calculated for a series of 1-minute bins over the cross-section of 195 trading pauses in our sample. The left (right) part of the chart represents the pre-pause (post-pause) period, while the five-minute trading pause is denoted by the vertical red line.

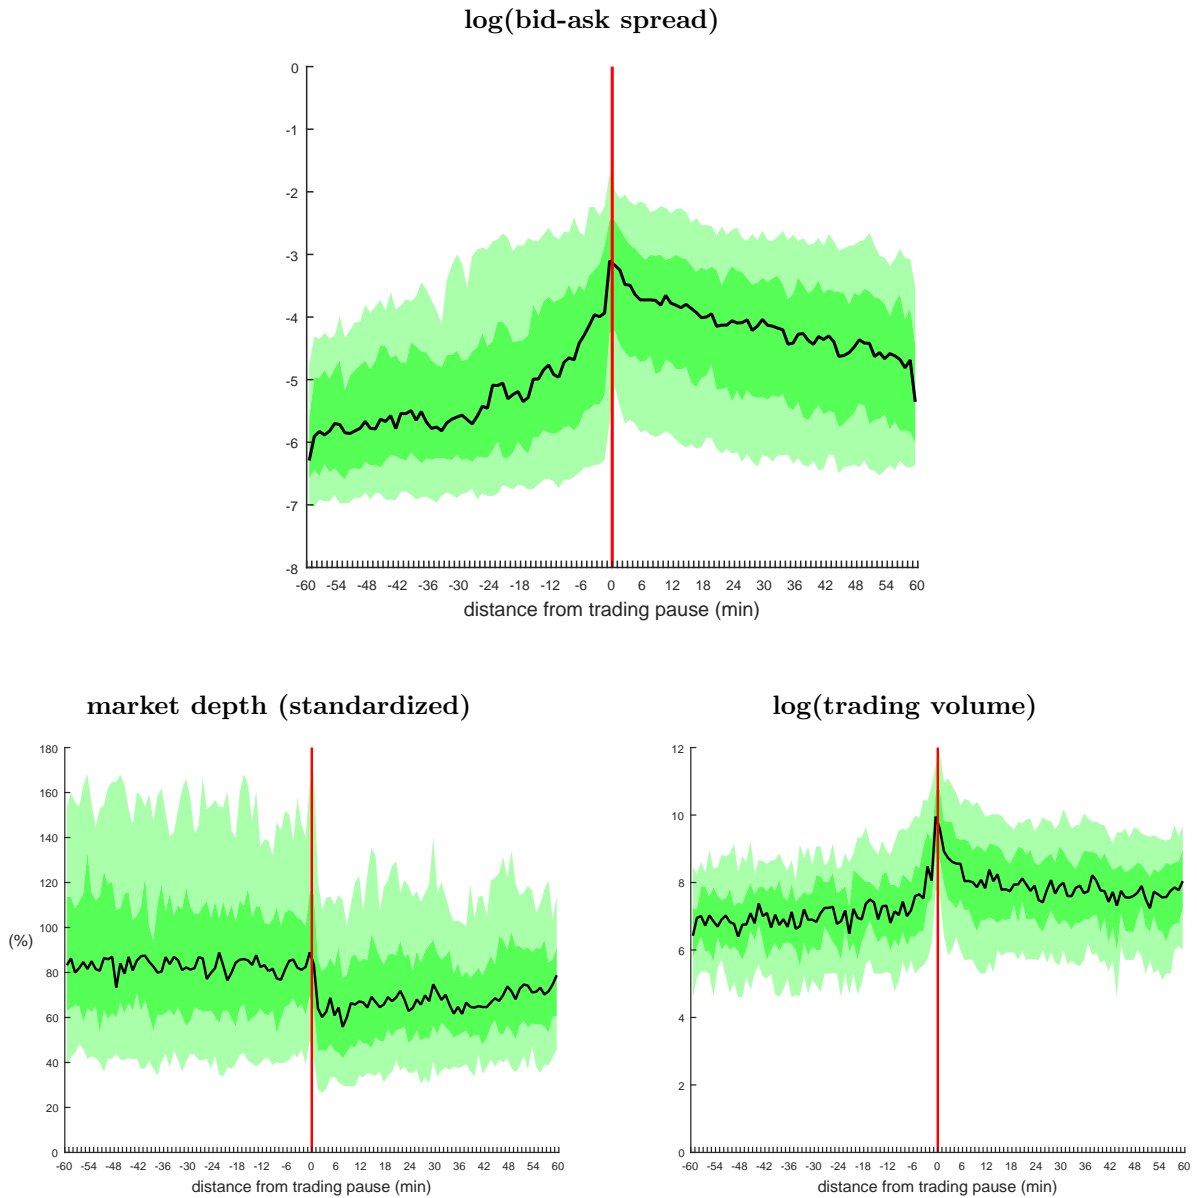


Figure 19: Market Asymmetry Measure Quantiles Around Trading Pauses

This figure shows how the cross-sectional quantiles of average (directional) order book imbalance (IMBAL) and (directional) cumulative market order flow (OFLOW) paths change in the two-hour interval around trading pauses. The plotted variables are defined in Equations 1 and 2 of Section 3.2. The shaded bands in the chart correspond to the 16th, 33rd, 50th (black line), 67th and 84th percentiles, calculated for a series of 3-minute bins over the cross-section of 195 trading pauses in our sample. The left (right) part of the chart represents the pre-pause (post-pause) period, while the five-minute trading pause is denoted by the vertical red line.

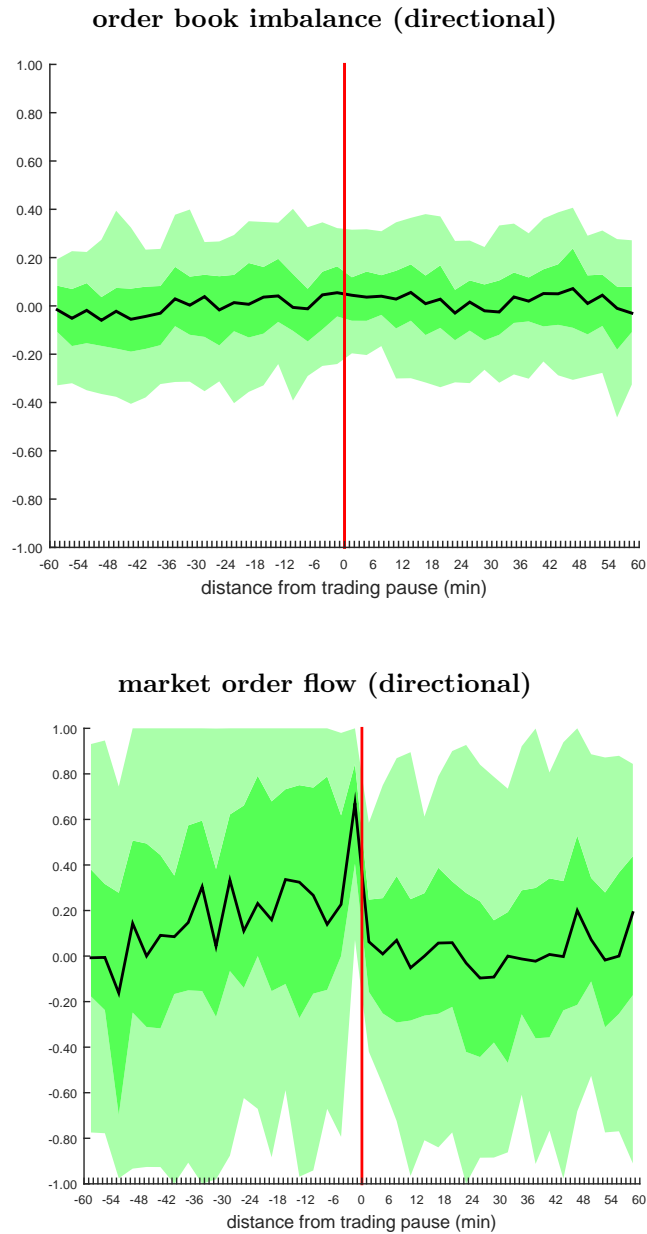
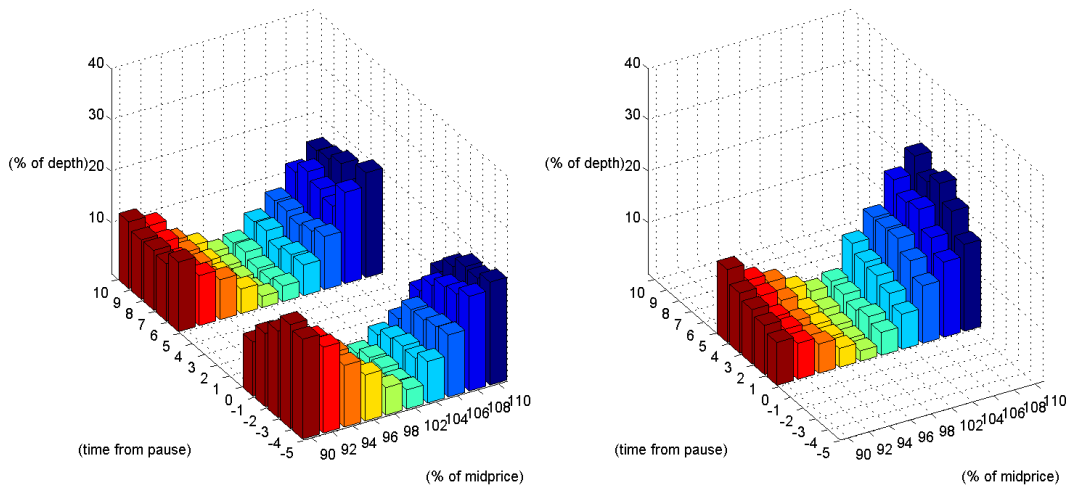


Figure 20: Median Order Book Depth Around Trading Pauses

This figure shows changes in the median standardized order book depth before (from -5 minutes to 0 minute on the y-axis), during (from 0 minute to 5 minutes on the y-axis) and after (from 5 minutes to 10 minutes on the y-axis) the 195 trading pauses in our sample. Around each trading pause, limit order book depth is measured every minute over a price grid (x-axis) centered around the mid-quote, where red (blue) is the bid (ask) side of the market. For the sake of robustness, the measured order book depth is standardized across stocks, and divided by the average order book depth of the respective stock in the month of the trading pause. The order book depth around upward (downward) trading pauses is shown in the upper (lower) panel. For the sake of better visibility, the five-minute trading pause period is plotted separately on the right hand side of each chart.

upward trading pauses



downward trading pauses

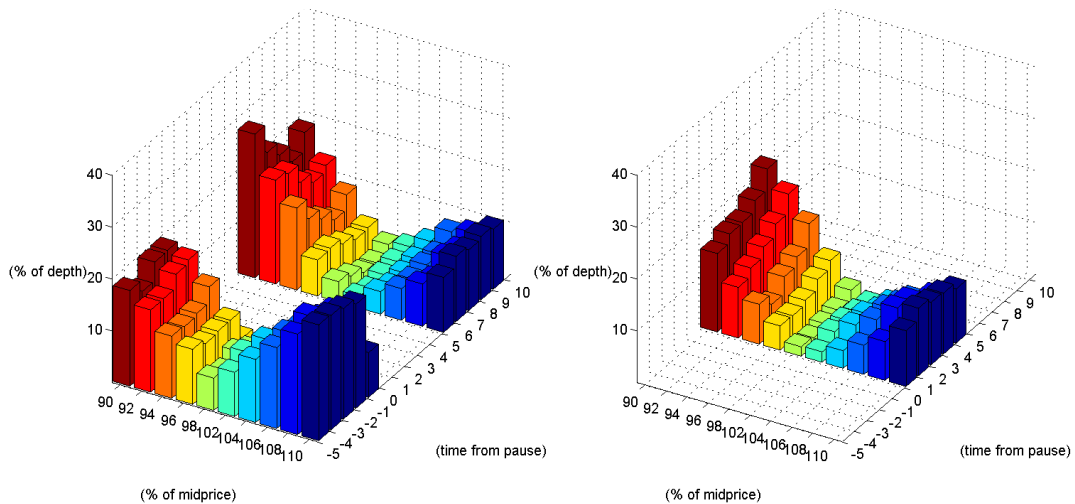
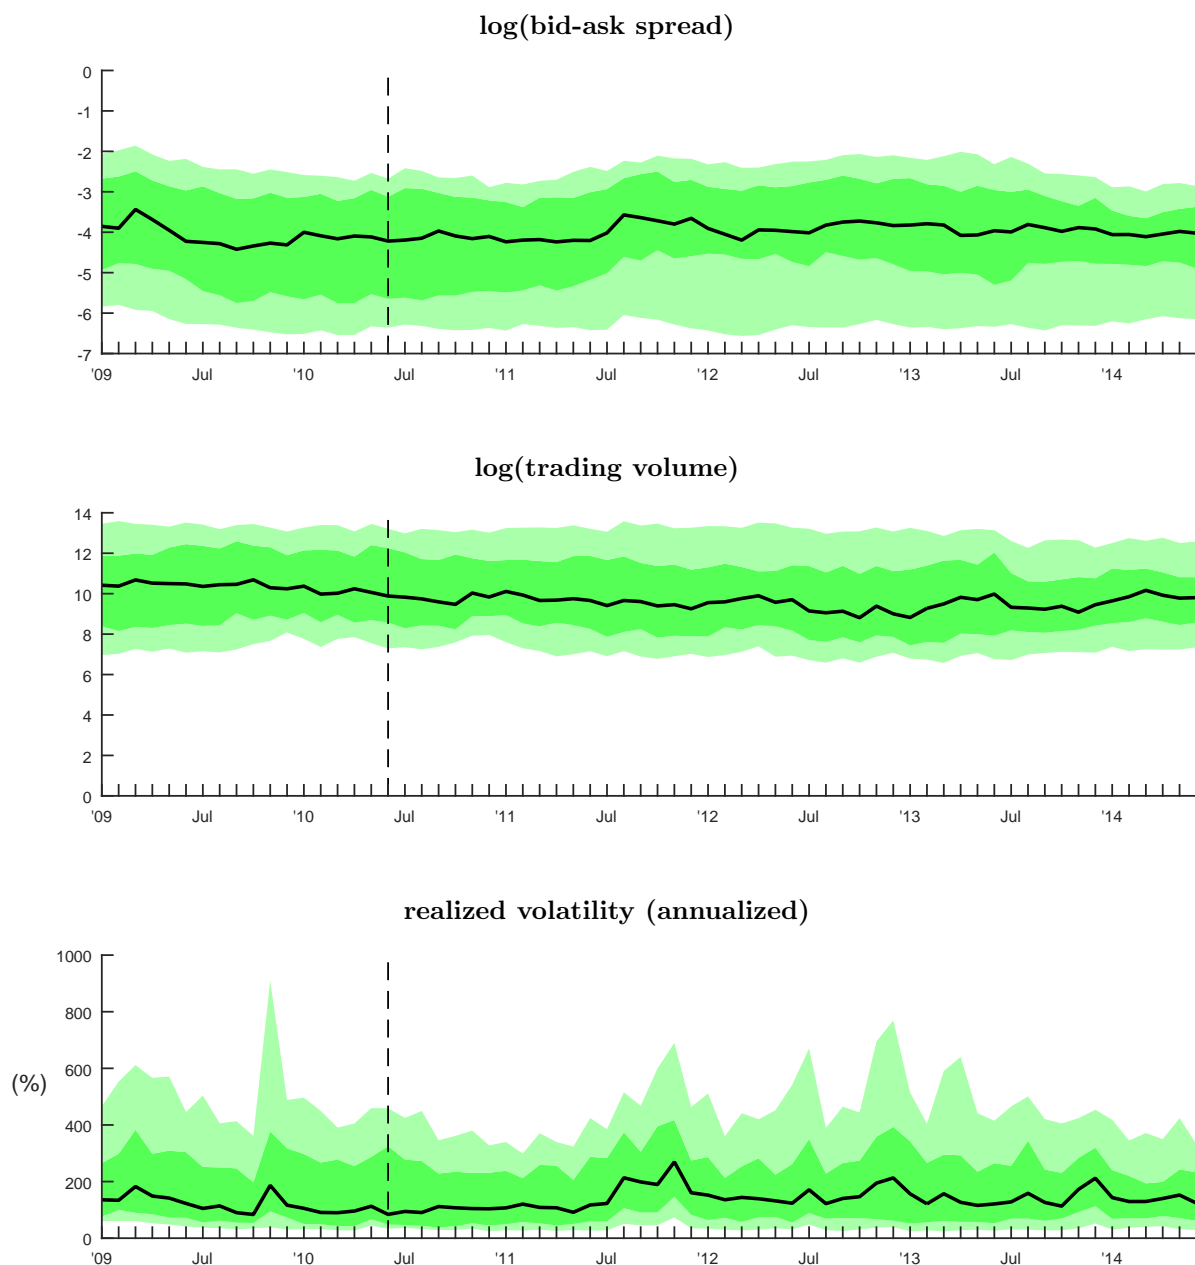


Figure 21: Comparing Market Measures Between the Control and Sample Periods

This figure shows how the distribution of (log) bid-ask spreads, (log) trading volume and (annualized) realized volatility changed over time, with special focus on the difference between the control (from January 2009 to May 2010) and sample (from June 2010 to June 2014) periods. The shaded bands in the chart correspond to the 16th, 33rd, 50th (black line), 67th and 84th percentiles, calculated on a daily level over the cross-section of 195 trading pauses in our sample. The left part of the chart represents the pre-regulation, the right part represents the post-regulation period, separated by the vertical dashed line.



Recent Issues

All CFS Working Papers are available at www.ifk-cfs.de.

No.	Authors	Title
570	Torben G. Andersen, Gökhan Cebiroglu, Nikolaus Hautsch	<i>Volatility, Information Feedback and Market Microstructure Noise: A Tale of Two Regimes</i>
569	Torben G. Andersen, Gökhan Cebiroglu, and Nikolaus Hautsch	<i>Volatility, Information Feedback and Market Microstructure Noise: A Tale of Two Regimes</i>
568	Joshua R. Goldstein, Christos Koulovatianos, Jian Li, Carsten Schröder	<i>Evaluating How Child Allowances and Daycare Subsidies Affect Fertility</i>
567	Roman Kräussl, Thorsten Lehnert, Denitsa Stefanova	<i>The European Sovereign Debt Crisis: What Have We Learned?</i>
566	Luiz Felix, Roman Kräussl, and Philip Stork	<i>Single stock call options as lottery tickets</i>
565	Luiz Felix, Roman Kräussl, and Philip Stork	<i>Implied Volatility Sentiment: A Tale of Two Tails</i>
564	Roman Kräussl and Elizaveta Mirgorodskaya	<i>The Winner's Curse on Art Markets</i>
563	Christiane Baumeister, Reinhard Ellwanger and Lutz Kilian	<i>Did the Renewable Fuel Standard Shift Market Expectations of the Price of Ethanol?</i>
562	Thierry Mayer, Marc Melitz and Gianmarco Ottaviano	<i>Product Mix and Firm Productivity Responses to Trade Competition</i>