A Preliminary Model for Estimating the Impact of Price Limits on Taiwan Stock Exchange[1]

Jeffery Russell, Chen Yang

Graduate School of Business

University of Chicago

December 17, 2003

[1]Research support from the Sanford J. Grossman Fellowship in Honor of Arnold Zellner is gratefully acknowledged; any opinions expressed herein are the authors' and not necessarily those of Sanford J. Grossman or Arnold Zellner.

Abstract

In many emerging stock markets, price limits are imposed on the magnitude of daily price movements. Price limit advocates claim that such limits serve as "circuit breakers" and decrease stock price volatility. Critics argue that the limits cause supply and demand imbalances in trading. Consequently, they prevent immediate corrections in price and increase the volatility of the opening return on the subsequent trading day. This is known as the volatility spillover hypothesis.

The empirical results on the impact of price limits in the stock market appear mixed. Furthermore, research using high frequency data directly from an emerging stock market where daily price limits are implemented is rare. Our study uses 5-minute intra-day price data from the Tai Wan stock exchange, where a daily price limit of 7% is imposed for all traded stocks. We consider the spillover effect of price limits as the effect that, given the price hitting the limits and staying at the limits for a certain amount of time, how much the next day's first five-minute(ten-minute, fifteen-minute...) volatility is elevated. We find significant evidence for this effect, after controlling for the company's specific overnight effect. Contrary to what might be expected after the emergence of a consensus price, the period immediately after a trading cessation is characterized by higher levels of volatility. We conclude that the daily price limits in TSE are ineffective in preventing over-reaction.

Contents

Introduction

In many emerging stock markets, price limits are imposed on the magnitude of daily price movements. Examples include the markets in Austria, Belgium, China, Japan, Korea, Malaysia, Mexico, Taiwan, and Thailand. This trading mechanism stands in sharp contrast to the primary practice in the U.S. stock market. The latter has obligations to provide continuous liquidity to the market, while the former restricts price movements in a certain range that is proportional to the previous day's closing price.

Price limit advocates claim that such limits serve as "circuit breakers" and decrease stock price volatility. They also believe that price limits will provide a cooling off period for rational reassessment during times of panic trading. On the other hand, critics argue that the limits cause supply and demand imbalances in trading. Consequently, they prevent immediate corrections in price and increase the volatility on the subsequent trading day. It is known as the volatility spillover hypothesis.

A large body of theoretical research reveals that the market microstructure is an important determinant of the components of volatility. Greenwald and Stein (1988) have argued that price limits could be beneficial if the price changes are caused by market failure, or a noise-generated panic. On the other hand, Fama (1989) points out that inciting trading in anticipation of circuit breakers will generally increase price volatility and merely delays the adjustment of prices to changes in fundamental values. In a study of trading halts in financial markets, Subrahmanyam (1994) shows that the existence of temporary circuit breakers may distort optimal trading decisions. Further, it leads to perverse effects on price variability and market liquidity. He concludes that a continuous price limits imposed daily will have the same volatility spillover effect as that of the temporary circuit breakers.

The empirical results on the impact of price limits in the stock market appear mixed. In cross-country research, Roll (1988) argues that price limits had no significant impact on the decline during the international crash of October 1987. On the contrary, Bertero and Mayer (1990) using the same data as Roll, find that price limits were effective. In studies that focus on single stock markets, Kuhn, Kurserk, and Locke (1991) show that stock volatility was not moderated by circuit breakers during the 1989 U.S. mini-crash. Kim and Rhee(1997) use Tokyo Stock Exchange data to conclude that volatility does not return to the normal level after reaching the price limits. They use daily data and compare the behavior of stocks that reach a price limit to stocks that almost reach their daily limit. However, Lauterbach and Ben-Zion(1993) study the behavior of Israeli stock market, and find price limits slightly smoothed return fluctuations during the October 1987 crash.

It is important to note that the methodology adopted by most of the above papers is very different from ours. They carry out event study and divide trading days into periods, namely, before the event and after the event. Then, different return windows are constructed and the null hypothesis that volatility is the same across all periods is tested. However, the event windows immediately following trading halts are necessarily high volatility periods, potentially inducing an unwanted selection bias. In our study, we condition on the volatility level and ask whether the volatility subsequent to a market closure is higher or lower than the expected given the volatility level immediatly preceding the close.

Though circuit breakers and trading halts have been studied in some markets that don't have daily price restrictions, like the New York Stock Exchange, empirical research using data directly from an emerging stock market where daily price limits are implemented is rare. In this paper, we use 5-minute intra-day return series from the Taiwan Stock Exchange, which imposes a 7% daily price limit for all the trading stocks.

The paper is organized as follows. Chapter 1 introduces the hypothesis of interest and its implication on policy issues. The methodology we use to test the hypothesis is also discussed. Chapter 2 provides a brief description of the TSE data. In Chapter 3 we specify an econometric model for the 5-minute return series and use this model to obtain the expected price path during halt. Results of the empirical tests are presented in Chapter 4. Finally we have a few comments about the model, and conclude the paper in Chapter 5. We find that, contrary to what might be expected after the emergence of a consensus price, the opening of the market after a trading cessation is characterized by higher levels of volatility.

Chapter 1

Hypothesis and Methodology

1.1 Motivation

Let P_t denote the price at time t, $t = 1, 2, ...T$. Let R_t denote the log return $ln(P_t) - ln(P_{t-1})$. Let $\mathcal{F}_t$ denote an information set available at time t which includes the filtration of returns and potentially other information available at time t. Let $\mathcal{I}_t$ be a random variable taking the value 1 if a trading halt is in effect at time t. $\mathcal{I}_t$ is not included in the information set $\mathcal{F}_t$. Let k denote the duration that price stays at the limits, our hypothesis of interest is then constructed by examining the following equality:

$$E[(R_{t,t+k} - \mu)^2 | \mathcal{F}_t, \mathcal{I}_t] = E[(R_{t,t+k} - \mu)^2 | \mathcal{F}_t] \tag{1.1}$$

Satisfaction of this equality implies that trading halts have no impact on volatility. Violations of this equality imply either volatility exuberance or cooling off effect, depending on whether the equality is replaced with a greater than or less than sign respectively. Of course we do not get to observe the "return" when a trading halt is in effect which complicates the testing of this equality.

To solve this unobservability problem we take advantage of the structure of the Tai Wan Stock Exchange(TSE). TSE imposes a daily price limit of 7% for all traded stocks. Within a trading day, the price for a single stock cannot move more than 7% from the previous day's closing price. During the sample period from January 3, 1998 to March 20, 1999, our summary statistics show that the price reached the limits, on average, 23.7 days out of the 324 trading days, which is about 7.3% of the sample period; and closed at the limits 17 days. In addition, the price stayed at the ceiling on average for 51 minutes once it reached the ceiling and it stayed at the floor for 44 minutes once it reached the floor. Thus within a trading day (3 hours), the stocks tend to stay at the limits for quite a long time once they reach the limits. Though stock price may remain on the limits for more than one day, this only occurs 9% of the time and for the rest, we are able to observe price movements when the market opens following a trading halt. More details of the data will be revealed in Chapter 2. With the above conditions, we can proceed to the next step of testing.

Right after the event day (on which the trading process was halted by price reaching the limit), the second day witnesses a free price movement when the market opens, due to reset of the price boundary. It is not clear what time interval in that day is the best to look at and connect to the impact of price limits, especially when we consider the policy implications that can be drawn from this study, whether elevation of volatility in the first 5-minute, 15-minute, 30-minute or 1-hour, etc... returns should be flagged to the policy makers. However, the test of "spillover effect" is indeed a test of the logic behind policy making. Moreover, it is expected to reveal, in a simple setting, the mechanism through which certain missing information during the halt is to affect the trading process and spill over to the next day. In this paper, we are going to perform tests using the first 5-minute, 15-minute, 30-minute, and 1-hour returns at the beginning of a post-event day, and document relative changes of the price limit effects.

Take the first 5-minute opening return as an example. Our basic assumptions are: equation (1.1) holds, and price can move freely and adjust immediately to the equilibrium level in the first five minutes. Under those two assumptions, the first 5-minute return is $lnP_{t^*+k+1} - lnP_{t^*}$, where t^* is the time that price reaches its limit on the previous day, and k is the duration of trading cessation. Furthermore, this difference should equal to the sum of all the expected returns on the unobservable price path over k. For simplicity, assume that the returns are uncorrelated, then a relationship between the observed and the hypothetical volatility in the first 5 minutes can be established:

$$R^{obs}_{t^*+k+1} = ln(P_{t^*+k+1}) - ln(P_{t^*}) = \sum_{j=1}^{k} R^{hyp}_{t^*+j} + R^{hyp}_{t^*+k+1} \tag{1.2}$$

$$Var(R^{obs}_{t^*+k+1}) = \sum_{j=1}^{k} Var(R^{hyp}_{t^*+j}) + Var(R^{hyp}_{t^*+k+1}) \tag{1.3}$$

Due to the summation term $\sum_{j=1}^{k} Var(R^{hyp}_{t^*+j})$ in equation (1.3)[1], volatility of the observed opening return will be clearly higher than that of the hypothetical one. This is the intuition behind the *Spillover Hypothesis*. The volatility in the post-halt period is elevated no matter how the market responds to the limits. Since stocks always experience price continuations, the same arguments apply to the volatility of the first 15-minute, 30-minute and 1-hour returns as well.

A simple test for a particular stock can be performed according to equation (1.3):

$$Var(R^{obs}_{t^*+k+1}) = \beta_1 Var(R^{hyp}_{t^*+k+1}) + \beta_2 * \sum_{j=1}^{k} Var(R^{hyp}_{t^*+j}) + \epsilon \tag{1.4}$$

[1] $R^{hyp}_{t^*+k+1}$ is the hypothetical opening return when the price limit restrictions were removed.

$$(R^{obs}_{t^*+k+1})^2 = \beta_1 E[(R^{hyp}_{t^*+k+1})^2|\mathcal{F}_{t^*}] + \beta_2 * \sum_{j=1}^{k} E[(R^{hyp}_{t^*+j})^2|\mathcal{F}_{t^*}] + \epsilon \tag{1.5}$$

Where $Var(R^{obs}_{t^*+k+1}) = (R^{obs}_{t^*+k+1})^2$, the realized volatility of the observed first 5-minute return; and $\sum_{j=1}^{k} Var(R^{hyp}_{t^*+j}) = \sum_{j=1}^{k}[E(R^{hyp}_{t^*+j})^2|\mathcal{F}_{t^*}]$, the sum of expected volatility over the halt period, which is calculated from GARCH forecasting and simulation. We apply a AR(3)-GARCH(2,2) model to describe the characteristics of a stock's 5-minute return series, this functional form takes into account the fact that real stock price has increased/decreased to the limit's level. We will discuss the modeling part in detail later. $Var(R^{hyp}_{t^*+k+1})$ is the volatility of the hypothetical first 5-minute return, in literature, it is also called "overnight volatility". For a particular stock, we assume the volatility of the first 5-minute return can be represented by a linear function of the realized volatility of the previous day. Under the null hypothesis that price limits have no effect, we estimate the linear function using data from all days that satisfy two conditions: (1) stock price does not reach the limit at the end of a day, (2) the first 5-minute return in the subsequent day is observable. Further, we apply the function to days that the stock price reaches the limit and causes a halt, in order to obtain the volatility of the expected hypothetical first 5-minute returns. More specifically, $E(R^{day_i}_{first-5})^2 = \omega_1 + \omega_2 * \sum_{all}^{day_{i-1}} R_t^2$, and $E[(R^{hyp}_{t^*+k+1})^2|\mathcal{F}_{t^*}] = \omega_1 + \omega_2 * \sum_{t^*-35+k}^{t^*} R_t^2$. Equation (1.3) reveals a mechanical impact of imposing price limits. However, we are more interested in testing whether there is derivative "spillover effect" besides the mechanical effect. Therefore, under the null hypothesis that equation (1.1) holds – the price limits do not have derivative "spillover effect", the βs in equation (1.5) should be one. In other words, the unobserved price path above the limits should not have a volatility leverage larger than one under the null. However, if the impact of price limits truly "spills over" to the next day more than can be explained by the mechanism shown in equation (1.2) and (1.3), we are going to see βs greater than one.

There can be many explanations if the leverage is greater than one. Suppose there is information diffusion coming along with a sudden increase or decrease in the price such that the price hits the limit. Imposing a boundary here merely delays the diffusion process as well as adding to uncertainty that people feel about the amount of information being released. The most important channel open to traders to judge the market's trend has been shut down, and price movement becomes invisible to investors. As a result, larger changes in prices, and consequently higher volatility, is supposed to compensate for the uncertainties caused by price limits.

On the other hand, it is also possible that the "abnormal" moves in price are initiated purely by noise traders. Instead of taking time to reassess their decision during the "cooling off" period, noise traders may become even more optimistic or pessimistic about the current situation, and push the price further away from the "normal" level. Noise trader effect, coupled with overnight effect, worsen the problem, and provide room for speculation on volatility dynamics. If that is the case, we expect to find a notably higher volatility than what is predicted by the mechanism in equation (1.3)– the "spillover" in the short run, but a reverse of price volatility in the long term. If we can group companies by criteria that noise traders use to make investment decisions, we may be able to find that certain companies have larger leverage than the others.

1.2 Random Effect Model

We ran a cross sectional random effect regression model to test the null hypothesis. The simplest case where a random effects model might be considered is when there are two stages of sampling. At the first stage units are selected at random from a population, and at the second stage a number of measurements is made sequentially in time or space on each unit sampled in the first stage. If we regard the stocks in our data set as fixed, then the most general model we might consider would have separate simple linear regressions of $Var(R^{obs}_{t^*+k+1})$ on $Var(R^{hyp}_{t^*+k+1})$ and $\sum_{j=1}^{k} Var(R^{hyp}_{t^*+j})$ for each stock, as shown in equation (1.4). However, there are large variations in slopes. We test if the slopes can be considered constant, as then we can adopt the simpler parallel-line regression model. The results suggest that differences between slopes are not adequately explained by a cross sectional regression with those parameters β_1 and β_2 fixed. We will discuss the results in detail in Chapter 4.

A promising way of generalizing the model is to assume that the stocks we examine are from a population where the slope coefficients depend on the sample. The model we investigate has the form

$$Y_{ij} = (\beta_1 + \theta_i^1)X_{ij}^1 + (\beta_2 + \theta_i^2)X_{ij}^2 + \epsilon_{ij} \tag{1.6}$$

where i denotes the sample and j the observation on that sample, and $\theta_i^1 \sim N(0, \sigma_1^2)$, $\theta_i^2 \sim N(0, \sigma_2^2)$, $\epsilon_{ij} \sim N(0, \sigma^2)$. Further, (θ_i^1, θ_i^2) and ϵ_{ij} are independent, but θ_i^1 and θ_i^2 can be correlated. This allows for the possibility that predictions from the fitted equation will need to cope with two sources of error, one associated with the sampling process and the other with the measurement process within the sample itself.

In particular, the parameters in the variance structure (such as the variance components) are estimated by maximizing the marginal likelihood of the residuals from a least-squares fit of the linear model, and then the fixed part of the effects are estimated by maximum likelihood assuming that the variance structure is known, which amounts to fitting by generalized least squares.

Suppose that for each company we write $Y_i = X_i\beta_i + \epsilon_i$, where $\beta_i = \beta + \lambda_i$,

$\lambda = \lambda(e, \theta)$, $E[\lambda_i] = 0$ and $E[\lambda_i \lambda_i'] = \Sigma$. Combining terms, we obtain the model

$$Y_i = X_i\beta + (\epsilon_i + X_i\lambda_i) = X_i\beta + \omega_i \tag{1.7}$$

where $E[\omega_i] = 0$ and $E[\omega_i\omega_i'] = \sigma_i^2 I + X_i\Sigma X_i' = \Pi_i$. For the full sample of observations from all companies, $V = [\Pi]$ with Π_i on the diagonal and zero off the diagonal in matrix $[\Pi]$.

$$V = [\Pi] = \begin{bmatrix} \Pi_1 & 0 & 0 & . & . & . & 0 \\ 0 & \Pi_2 & 0 & . & . & . & 0 \\ & & . & & & & \\ & & . & & & & \\ & & . & & & & \\ 0 & 0 & 0 & . & . & . & \Pi_n \end{bmatrix}$$

To estimate the parameters, let b_i be the i^{th} ordinary least squares coefficient estimator for the i^{th} company,

$$Var[b_i|X_i] = E[(b_i - \beta_i)(b_i - \beta_i)'|X_i] = E[(X_i'X_i)^{-1}X_i'\omega_i\omega_i'X_i(X_i'X_i)^{-1}|X_i] \tag{1.8}$$

and since $E(\omega_i\omega_i') = \Pi_i$, let

$$V_i + \Sigma = (X_i'X_i)^{-1}X_i'\Pi_i X_i(X_i'X_i)^{-1} \tag{1.9}$$

be the covariance matrix of b_i, where $V_i = \sigma_i^2(X_i'X_i)^{-1}$. On the other hand, the covariance matrix of $\hat{\beta}$, the GLS estimator, is $Var[\hat{\beta}_i] = [X_i'\Pi_i^{-1}X_i]^{-1}$.

Rao (1973) proved that the GLS estimator in the cross-sectional heteroscedastic model may be written as:

$$\hat{\beta} = \sum_{i=1}^{n} \left[\left(\sum_{i=1}^{n} (V_i + \Sigma)^{-1} \right)^{-1} (V_i + \Sigma)^{-1} \right] b_i \tag{1.10}$$

indicating that GLS estimator is a matrix weighted average of the OLS estimators. In order to estimate the unknown parameters in V_i and Σ, Swamy (1971) suggested a two-step approach. First we obtain the ordinary least squares estimators b_i for each company, and get the sample covariance matrix $\hat{V}_i = s_i^2(X_i'X_i)^{-1}$, where $s_i^2 = e_i'e_i/(n_i - k)$, as we have k parameters and n_i observations in each company.

Then the second step is to let $\bar{b} = \frac{1}{n}\sum_{i=1}^{n} b_i$.

$$\hat{\Sigma} = E[\lambda_i \lambda_i'] = \frac{1}{n-1}\left(\sum_{i=1}^{n} b_i b_i' - n\bar{b}\bar{b}'\right) - \frac{1}{n}\sum_{i=1}^{n} \hat{V}_i \qquad (1.11)$$

This comes from equation (1.8). Intuitively, there should be an adjusting component $\frac{1}{n}\sum_{i=1}^{n} \hat{V}_i$, because the between-group variation among b_i's calculated from the n least square equations does not control for the within-group variation among the observations of X_i's. So to calculate an estimator for Σ in the cross-sectional model, we need to make the above adjustment. In large-sample behavior of $\hat{\Sigma}$, the second matrix will be negligible.

To get the estimators of variance components, we use maximum likelihood estimation.

$$lnL = -\frac{n}{2}ln(2\pi) - \frac{1}{2}ln\,|\Pi| - \frac{1}{2}\omega'(\Pi)^{-1}\omega \qquad (1.12)$$

where $\Pi = \Pi(V_i, \Sigma)$. Let $\Gamma = \Pi^{-1}$, and use $Y - X\beta$ to substitute ω,

$$lnL = -\frac{n}{2}ln(2\pi) + \frac{1}{2}ln\,|\Gamma| - \frac{1}{2}(Y - X\beta)'\Gamma(Y - X\beta) \qquad (1.13)$$

Then

$$\frac{\partial lnL}{\partial \beta} = X'\Gamma(Y - X\beta) \qquad (1.14)$$

$$\begin{aligned} \frac{\partial lnL}{\partial \Gamma} &= \frac{1}{2}\left[\Gamma^{-1} - \omega\omega'\right] \\ &= \frac{1}{2}\left[\Pi - \omega\omega'\right] \end{aligned} \qquad (1.15)$$

Given a consistent estimator of $\Pi(V_i, \Sigma)$ from equation (1.10), we use (1.13) above to estimate β. With the new estimates of β, we then reestimate all the parameters in Π or Γ from (1.14). Thereafter we repeat the two-step iteration until satisfactory convergence has been achieved.

Denoting $A_i = (\Sigma^{-1} + V_i^{-1})\Sigma^{-1}$, $\Omega_i = [\sum_{i=1}^{n}(V_i + \Sigma)^{-1}]^{-1}(V_i + \Sigma)^{-1}$ and $V_i = \sigma_i^2(X_i'X_i)^{-1}$,

$$\begin{aligned} \hat{\beta}_i &= [\Sigma^{-1} + V_i^{-1}]^{-1}[\Sigma^{-1}\hat{\beta} + V_i^{-1}b_i] \\ &= A_i\hat{\beta} + [I - A_i]b_i \end{aligned} \tag{1.16}$$

For standard errors and confidence intervals,

$$Var[\hat{\beta}_i] = \begin{bmatrix} A_i \\ I - A_i \end{bmatrix} \begin{bmatrix} \sum_{i=1}^{n} \Omega_i(\Sigma + V_i)\Omega_i' & \Omega_i(\Sigma + V_i) \\ (\Sigma + V_i)\Omega_i' & (\Sigma + V_i) \end{bmatrix} \begin{bmatrix} A_i \\ I - A_i \end{bmatrix} \tag{1.17}$$

In $S+$, the test of the random coefficients model is based, equation by equation, on the differences between the OLS estimates and a weighted average of the OLS estimates. The test statistic suggested by Swamy (1971) is:

$$\chi^2 = \sum_{i=1}^{n} [b_i - \tilde{b}]'\hat{V}_i^{-1}[b_i - \tilde{b}] \tag{1.18}$$

where $\tilde{b} = \left[\sum_{i=1}^{n} \hat{V}_i^{-1}\right]^{-1} \sum_{i=1}^{n} \hat{V}_i^{-1}b_i$. The statistic is asymptotically distributed as chi-squared with K(n-1) degress of freedom under the null hypothesis of parameter constancy. It is equivalent to the standard F statistic for testing $H_0 : \beta_1 = \beta_2 = ... = \beta_n$ in the generalized model $Y_i = X_i\beta_i + \epsilon_i$ $(i = 1, 2, ...n)$ with $E[\epsilon_i\epsilon_i'] = \sigma_i^2 I$, for $i = j$.

Chapter 2

Presentation of Data

2.1 Overview

The Tai Wan Stock Exchange (TSE) is an order-driven call market that does not utilize designated market makers. Investors issue orders, and the market uses a periodic batch process to match demand and supply. The clearing price is determined in order to maximize the trading volume. Tai Wan Stock Exchange (TSE) imposes a daily price limit of 7% for all traded stocks. Within a trading day, the price for a single stock cannot move more than 7% from the previous day's closing price. Therefore, the maximum 1-day return is 7% and the minimum 1-day return is -7%.

The trading hours are from 9:00 am to 12:00 noon on weekdays. Trading also occurred on the first and third Saturdays of each month and from 9:00 am to 11:00 am until March 1998. From April 1998, the trading hours for these two Saturdays are extended to 3 hours from 9:00 am to 12:00 noon. The original dataset consists of 5-minute return series of 346 listed companies in the TSE from January 3, 1998 to March 20, 1999, all adjusted for dividends and splits. The TSE categorizes these companies into 18 industrial sectors.

Table 2.1: TSE category of 18 industrial sectors

Code	Sector	# Firms	A	B	C	D	E	F
11	Cement	7	15.0	9.4	3.0	37.0	43.9	82.1
12	Food	24	29.4	21.6	11.3	48.5	46.6	59.8
13	Plastics	16	29.4	22.7	7.4	51.0	46.1	227.9
14	Textiles	48	27.0	18.8	7.8	55.3	41.7	72.8
15	Machinery	19	26.6	19.3	8.5	52.0	54.6	50.8
16	Elec. App.	12	21.6	15.4	6.5	45.1	50.5	102.1
17	Chemicals	16	24.2	17.1	6.9	48.0	38.8	65.2
18	Glass	6	18.8	13.0	5.0	59.9	41.6	51.8
19	Paper Pulp	6	21.2	15.5	5.5	57.6	38.3	85.8
20	Iron Steel	25	33.0	24.0	14.3	53.6	49.7	83.5
21	Rubber	8	27.3	20.5	7.4	60.0	46.0	82.6
22	Automobile	4	6.8	4.3	1.0	57.6	47.0	70.7
23	Electronics	66	36.8	27.1	8.4	49.4	49.1	278.4
25	Construction	31	31.1	22.8	11.2	45.7	51.1	96.1
26	Shipping	14	18.6	13.2	4.1	49.6	28.0	83.1
27	Tourism	6	20.3	13.5	5.1	49.7	29.6	30.3
28	Bank	34	18.5	13.6	5.1	52.1	30.2	166.2
99	Other	4	20.8	13.8	5.8	52.0	29.5	85.6
	Tot/Ave	346	23.7	17.0	6.9	50.9	44.1	98.6

***A**: Sector average of days reaching the limits. **B**: Sector average of days closing at the limits. **C**: Sector average of consecutive days closing at the limits. **D**: Sector average duration of staying at the ceiling. **E**: Sector average duration of staying at the floor. **F**: Sector median of daily trading volume.*

Table 1 gives summary statistics of how often the price limits are reached for the firms classified by industrial sectors. During the sample period from January 3, 1998 to March 20, 1999, the price reached the limits, on average, 23.7 days out of the 324 trading days, which is about 7.3% of the sample period. The stocks closed at the limits on average 17 days out of the 324 trading days. This implies that once the limits are reached in a trading day, it is more likely to close at the limits. Table 1 also summarizes how long the price stays at the limits. On average, the price stays at the ceiling for 51 minutes once it reaches the ceiling and it stays at the floor for 44 minutes once it reaches the floor. Thus within a trading day (3 hours), the stocks tend to stay at the limits for quite a long time after they reach the limits. It is also noticed that sector 23–Electronics, has the highest median daily trading volume. The most heavily traded company in Tai Wan, TSMC (Taiwan Semiconductor Manufacturing Company Ltd.), is included in this sector. On the other hand, sector 27–Tourism, has the lowest median daily trading volume. Table 1 clearly shows that the daily price limits play a very important role during the sample period in the Tai Wan Stock Exchange. Therefore, the data from TSE is a natural candidate for us to study the effect and impact of the daily price limits.

2.2 Typical Firms

For illustration purpose, one firm from each category is picked out (18 firms in all) and their data structure is analyzed. There are 36 observations for a typical trading day, and 24 for each special Saturday as explained in the previous section. The variable *Overnight Duration* is constructed to summarize the length of time that stock price stays at the limits — and hence no transaction occurs — before the market opens and resumes to trade on the next day. By implication, *Overnight Duration>0.*

Figure 1: Mean of Overnight Duration

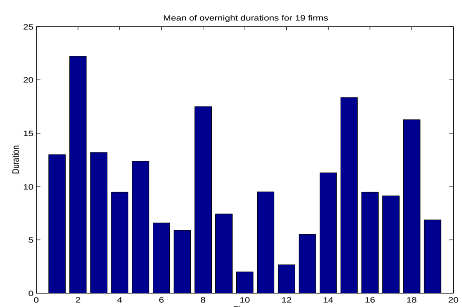

Figure 1 shows the mean of overnight durations for the 18 firms. The second firm, TAIWAN PINEAPPLE, in sector 2 (Food) has the largest overnight duration. For this stock, if the market closes at its limit on one day and then resumes to trade on the following day, the price would have stayed at the limit for an average of 2 hours on the previous day.

Further examination on TAIWAN PINEAPPLE is presented in Figure 2. Graph(a) reveals a spike at *overnight duration=92* around day 277. Graph (b) on the right provides an explanation for this number. The price has a clear upward trend before day 277, and in fact, it has been staying at the 7% ceiling for two and a half consecutive days. In total, there are exactly 92 observations. Immediately after that day, the open-to-open price seems to be very volatile on the stock market. This observation is consistent with the prediction of the *spillover hypothesis*.

Figure 2: TAIWAN PINEAPPLE Price Data and Duration Plots

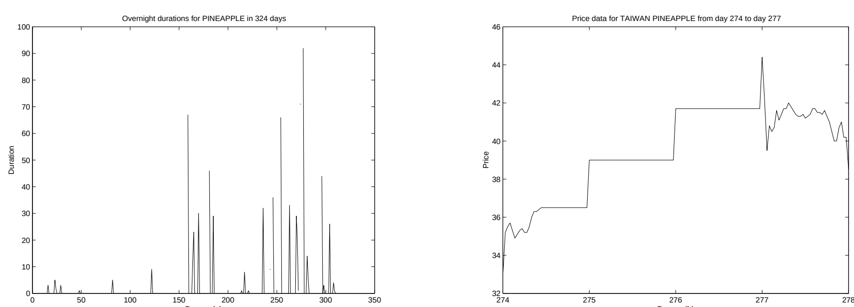

However, when performing empirical tests, we mainly focus on the cases where *overnight duration* is less than one day (amount to 36 observations). There are two reasons behind this. First, the GARCH model we use to forecast the price movement during the halt period is not good for prediction over long time horizon, especially across days. The TSE market opens three hours every day and closes 21 hours before the next opening. It is not desirable to make forecasts across to the next day when such a significant "overnight" period is present. Second, frquency study of all the stocks shows that *overnight duration* <= 36 accounts for 91% of the times, when *overnight duration* > 0 are counted. Therefore, we focus on the majority cases and include in our study only the first two days of the minority cases, in which the duration is longer than one day. To be more specific, for *overnight duration* > 36, the observed opening return after the first event day is |7%| (at the limit); we include this number in our study even though it underestimates the true absolute return at that time. As a result, we minimize the downward bias that may be caused by leaving those highly volatile, post-event days entirely out of the study.

The original TSE data that we obtained is 5-minute price series. The return is defined conventionally as $R_t = ln(P_t) - ln(P_{t-1})$, where t represents each 5-minute interval. Another restriction on TSE is that the price cannot move more than 2 tic sizes between two consecutive transactions during a day. In practice, trade is usually not observed to occur when the price moves beyond 6.8% or below −6.8% of the previous day's closing price.

2.3 Feature of Data

It is well known that the volatility is higher at the beginning of a trading day than during that day. Figure 3 plots the intra-day volatility pattern for the 346 TSE stocks in our dataset. The volatility is defined as the standard deviation of the 5-minute returns.

The TSE market opens at 9:00am and closes at 12:00 noon to make 36 5-minute time bins. For each stock, we calculate the standard deviation of the returns for each of the 36 time bins. Then the standard deviations are averaged over 346 stocks to create the average intra-day volatility pattern. The following graph presents a clear U-shape with the highest values achieved at two ends, i.e., the opening and closing

Figure 3: Average Intra-day Volatility Pattern

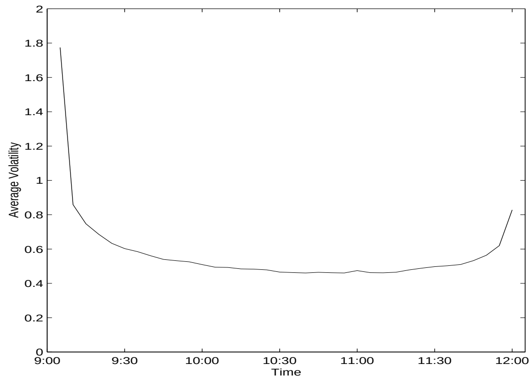

of the market. Therefore, we apply a widely used method–standardization to handle this deterministic volatility pattern.

Specifically, for a given stock, let $r_{t,k}$ be the 5-minute return at day t and k_{th} bin and

$$\delta_k = \sqrt{\frac{\sum_{i=1}^{324}(r_{i,k} - \bar{r}_k)^2}{323}}$$

be the sample standard deviation of the return, where $\bar{r}_k = \sum_{i=1}^{324} r_{i,k}/324$. The standardized 5-minute return of the stock is then defined as

$$RET_{t,k} = \frac{r_{t,k}}{\delta_k}.$$

Under such transformation, the U-shape pattern in the volatility has been removed.

Chapter 3

Econometric Models

Andersen and Bollerslev(1997) have shown that high frequency stock returns usually have heavy tails, and that their volatility is correlated. In convention, the generalized autoregressive conditional heteroskedasticity (GARCH) model is used to capture heavy tails and volatility clusterings. The GARCH(p,q) process models the conditional variance h_t as a linear function of past squared innovations and past conditional variances. It incorporates such a fact in the financial markets that large innovations tend to be followed by large innovations when uncertainty is prevalent.

In estimating the GARCH model for all the stocks, we exclude those returns when the relevant prices are at the limits, because once the price reaches its limit, it can stay at the limit or move only in one direction and hence exhibit unusual dynamics. Once the price departs from the limit, the next return is included in the analysis. Under the null hypothesis, the GARCH model estimated should govern the trading process if there were no daily price limits imposed on the market, i.e. price limits were not effective to change the volatility dynamics. As such, we will proceed to use the GARCH model and forecast the price movements where they were not observed.

We find the TSE data typically exhibit strong negative serial correlations up to the third lag. Therefore, we use AR(3) as the conditional mean function for the

GARCH model. The conditional mean equation is specified as follows:

$$R_t = \omega_0 + \omega_1 R_{t-1} + \omega_2 R_{t-2} + \omega_3 R_{t-3} + z_t \tag{3.1}$$

Where R_t is defined as $R_{t,k} = \frac{r_{t,k}}{\delta_k}$. The sample standard deviation of the return δ_k is defined as $\sqrt{\sum_{i=1}^{324}(r_{i,k} - \bar{r}_k)^2/323}$, where $\bar{r}_k$ is $\sum_{i=1}^{324} r_{i,k}/324$.

In this paper, a GARCH(2,2) process finally fits better than other alternatives, and the ranks also agree with the GARCH model identified in Cho, Russell, Tsay and Tiao (2002). This model is applied to the conditional variance of the returns,

$$h_t = \alpha_0 + \alpha_1 z_{t-1}^2 + \alpha_2 z_{t-2}^2 + \beta_1 h_{t-1} + \beta_2 h_{t-2} \tag{3.2}$$

Where $z_t = \sqrt{h_t}\varepsilon_t$. As in all GARCH models, $\alpha_i \geq 0$,$\beta_j \geq 0$, and $\sum_{i=1}(\alpha_i + \beta_i) < 1$. The unconditional variance of z_t is finite whereas its conditional variance h_t evolves over time. It is straightforward that

$$E(z_t^2) \; = \; \frac{\alpha_0}{1 - \sum_{i=1}(\alpha_i + \beta_i)} \tag{3.3}$$

Since we have standardized the volatility to remove the U-shape pattern, we expect $\alpha_0 \; = \; 1 - \sum_{i=1}(\alpha_i + \beta_i)$.

The model we want to test consists mainly of three parts:

$$Vol_{ij}^{obs} = (\beta_1 + \theta_i^1) * Vol_{ij}^{hyp} + (\beta_2 + \theta_i^2) * Vol_{halt}^{hyp} + \epsilon_{ij} \tag{3.4}$$

On the left hand side, the volatility of the observed first 5-minute return is approximated by realized volatility, i.e., square of the return in its original units. On the right hand side, the hypothetical volatility of the first 5-minute return is estimated by the linear function specific to each firm of the realized volatility from the previous day. The hypothetical volatility over the halt period will be forecasted by the AR(3)-GARCH(2,2) model. Since we know the unobserved price paths during the halt can only be above the 7% limits or below the −7% limits, the forecasts are conditional ones. Accordingly, for each case that the price reaches the limit and stays there until the end of the day, we simulate from the AR(3)-GARCH(2,2) model 1000 paths that are all above or below the limits. After that, 1000 returns for each entire path are calculated and volatility of the 1000 returns is recorded. Because the returns used to build the GARCH model are standardized, the volatility forecast is not on the same scale as the realized volatility. Thus, we "standardize back" the volatility forecast to its original unit. We use days that have *overnight duration* > 0, because only those observations are relevant to the null hypothesis we want to test.

In order to implement the random effect model, multiplicates from each company are required. However, there are 14 companies that have single pair of observations (observed volatility and expepcted volatility over the halt), and 21 companies that have only two pairs of observations over the sampling period. Hence, these companies are eliminated from the dataset. Since we are looking at a narrow window of the first five-minute return of each post-event day, stocks that in general are illiquid and trade very little in that five minutes may behave undesirably and contaminate our results. By defining an illiquid stock as one that is traded less than five times on average in the first five minutes, we find that nine companies are "illiquid", and they are excluded as well. Furthermore, due to uncontrolled factors in the numerical estimation of the GARCH model, some coefficients obtained for the GARCH coefficients contribute to forecasting negative volatilities, which are not feasible to proceed with. As a result, the involved compnanies (11 in total) are taken out. Therefore, in the end we have 291 companies and 2410 pair of observations to run the analysis.

Chapter 4

Estimation Results

We use S+ to estimate the coefficients of the *Random Effect Model*. t^* is the time when the price hits the limit.

$$[R^{obs(ij)}_{t^*+k+1}]^2 = (\bar{\beta}_1 + \theta^1_i) * E[(R^{hyp(ij)}_{t^8+k+1})^2|\mathcal{F}_{t^*}] + (\bar{\beta}_2 + \theta^2_i) * E[(R^{hyp(ij)}_{t^*,t^*+k})^2|\mathcal{F}_{t^*}] + \epsilon_{ij} \quad (4.1)$$

According to Swamy's test statistic in equation (1.18), we first show that the random effects are significant and cannot be explained by a fixed coefficient regression.

Table 4.1: Random Coefficient VS. Fixed Coefficient

Model	AIC	BIC	loglik	Test	L. Ratio	p-value
Random Coef.	19045.67	29083.50	-14516.84			
Fixed Coef.	29092.47	29117.69	-14542.24	1 vs 2	50.798	<0.0001

Second, we run a random effect model with a free intercept and fixed slopes: $[R^{obs(ij)}_{t^*+k+1}]^2 = \alpha + \eta_i + \beta_1 * E[(R^{hyp(ij)}_{t^*+k+1})^2|\mathcal{F}_{t^*}] + \beta_2 * E[(R^{hyp(ij)}_{t^*,t^*+k})^2|\mathcal{F}_{t^*}] + \epsilon_{ij}$. The results are shown in the next page.

Under the null hypothesis, the intercept should be zero, and the fixed slope coefficients should be one. However, from the table we find the intercept is significantly larger than zero, and the slope coefficients are larger than one. Third, we

Table 4.2: Random Effect Model by Comp. – Intercept

Random Part	StdDev
Intercept η_i	4.188
Residual ϵ_{ij}	12.052

Fixed Part	Value	Std. Error	t-value vs. 0	t-value vs. 1
Intercept α	5.856	0.838	6.988	
Slope β_1	1.666	0.307	5.421	2.169
Slope β_2	1.147	0.041	28.595	3.585

run a random effect model with both intercept and slopes being free: $[R^{obs(ij)}_{t^*+k+1}]^2 = \alpha + \eta_i + (\bar{\beta}_1 + \theta^1_i) * E[(R^{hyp(ij)}_{t^*+k+1})^2|\mathcal{F}_{t^*}] + (\bar{\beta}_2 + \theta^2_i) * E[(R^{hyp(ij)}_{t^*,t^*+k})^2|\mathcal{F}_{t^*}] + \epsilon_{ij}$.

Table 4.3: Random Effect Model by Comp. – Intercept and Slope

Random Part	StdDev
Intercept η_i	2.348
Slope θ^1_i	1.331
Slope θ^2_i	0.249

Fixed Part	Value	Std. Error	t-value vs 0	t-value vs 1
Intercept α	5.693	0.871	6.531	
Slope $\bar{\beta}_1$	1.705	0.351	4.869	2.009
Slope $\bar{\beta}_2$	1.173	0.047	24.955	3.681

Again, all the fixed parts of the slope coefficients are significantly larger than one. The intercept is greater than zero. It indicates for each observation, the total spillover effect on volatility is elevated by a fixed amount. Therefore, we reject the null hypothesis and conclude that the price limits have a significant spillover effect on the opening return of the next day, and such effect is not explained by the mechanical leverage revealed in equation (1.3).

Fourth, we group the companies by sector and examine whether there are significant discrepencies among sectors. Totally we have 17 sectors, s=1,2,...17, assuming each group is a random draw from the population. For model $[R^{obs(sj)}_{t^*+k+1}]^2 =$

$\alpha + \eta_s + \beta_1 * E[(R^{hyp(sj)}_{t^*+k+1})^2|\mathcal{F}_{t^*}] + \beta_2 * E[(R^{hyp(sj)}_{t^*,t^*+k})^2|\mathcal{F}_{t^*}] + \epsilon_{sj}$.

Table 4.4: Random Effect Model by Sector – Intercept

Random Part	StdDev
Intercept η_s	2.399
Residual ϵ_{sj}	12.431

Fixed Part	Value	Std. Error	t-value vs 0	t-value vs 1
Intercept α	6.308	0.986	6.398	
Slope β_1	1.704	0.288	5.918	2.444
Slope β_2	1.165	0.039	29.148	4.231

For model: $[R^{obs(sj)}_{t^*+k+1}]^2 = (\alpha + \eta_s) + (\bar{\beta}_1 + \theta^1_s) * E[(R^{hyp(sj)}_{t^*+k+1})^2|\mathcal{F}_{t^*}] + (\bar{\beta}_2 + \theta^2_s) * E[(R^{hyp(sj)}_{t^*,t^*+k})^2|\mathcal{F}_{t^*}] + \epsilon_{sj}$.

Table 4.5: Random Effect Model by Sector – Intercept and Slope

Random Part	StdDev
Intercept e_s	2.135
Slope θ^1_s	0.765
Slope θ^2_s	0.134

Fixed Part	Value	Std. Error	t-value vs 0	t-value vs 1
Intercept α	6.291	0.955	6.586	
Slope $\bar{\beta}_1$	1.647	0.367	4.482	1.763
Slope $\bar{\beta}_2$	1.201	0.057	21.207	3.526

We also apply the random effect model within each individual sector and do not find significant differences between sectors in terms of the spillover effect. In other words, the effect is robust and affects individual sector as well as company. There is no sector-specific effect involved. See Table 4.6.

Fifth, to further check robustness, we group the companies by liquidity levels. In this paper, the liquidity for each company is defined as the daily mean of dollar trading volumes (MDTV). We roughly separate all the companies into three quantiles: low liquidity (MDTV<25.78 million), mid liquidity (25.78<MDTV<79 million) and high liquidity (MDTV>79 million), and each category includes about 804 observations. See Table 4.7 and Table 4.8. t values against 1 instead of 0 are reported in all tables.

Table 4.6: Random Effect Model for Individual Sector

fix slope					random slope				
Sector	$\hat{\beta}_1$	std($\hat{\beta}_1$)	$\hat{\beta}_2$	std($\hat{\beta}_2$)	Sector	$\bar{\beta}_1$	θ_1	$\bar{\beta}_2$	θ_2
1	4.419	2.184	2.088	0.539	1	4.419	0.005	2.089	0.002
2	1.243	1.233	1.133	0.157	2	1.994	1.446	1.158	0.263
3	2.889	1.291	0.971	0.139	3	2.498	0.836	0.914	0.089
4	3.394	0.417	1.456	0.125	4	4.386	0.486	1.423	0.137
5	4.493	1.727	1.256	0.200	5	5.157	1.778	1.265	0.127
6	2.086	0.832	1.647	0.241	6	2.494	0.989	1.688	0.104
7	3.719	1.714	1.198	0.198	7	3.664	1.717	1.222	0.197
8	1.814	0.003	2.268	0.365	8	1.849	0.037	2.246	0.196
9	1.963	0.971	1.119	0.341	9	2.275	0.721	1.499	0.586
10	2.636	1.015	0.854	0.151	10	2.587	0.628	0.975	0.409
11	3.829	1.021	1.157	0.204	11	4.134	1.007	1.140	0.116
13	1.838	0.628	1.021	0.069	13	1.822	1.069	1.038	0.284
14	3.492	0.948	1.224	0.127	14	3.481	1.719	1.219	0.194
15	3.424	0.969	1.189	0.234	15	2.675	1.304	1.494	0.352
16	3.344	1.465	2.221	0.746	16	3.345	1.465	2.220	0.747
17	3.449	0.564	1.393	0.181	17	3.727	1.800	1.462	0.424
18	4.758	2.337	1.582	0.520	18	5.187	1.212	1.634	0.451

For stocks with a mean daily dollar trading volume smaller than 79 million, the slope coefficient of the forecasted volatility during the halt is significantly greater than one in both random effect models. For stocks with mean daily dollar trading volume larger than 79 million, the slope coefficient of the forecasted volatility during the halt is not different from one. On the other hand, the slope coefficient of the estimated first 5-minute return volatility is not different from one in all the regressions. Therefore, the spillover effect varies depending on the liquidity level of stocks, but there is no cooling off effect at any level identifed.

Sixth, as discussed in Chapter 1, we also want to test the spillover effect for the

Table 4.7: Random Effect Model by Liquidity. – Intercept

	low liquidity		mid liquidity		high liquidity	
Random Part	StdDev	#Firms	StdDev	#Firms	StdDev	#Firms
Intercept η_i	2.262	811	4.175	804	2.674	796
Residual ϵ_{ij}	12.958		12.356		10.655	
Fixed Part	Value	t-value-1	Value	t-value-1	Value	t-value-1
Intercept α	8.687		4.234		3.362	
Slope $\hat{\beta}_1$	1.395	0.859	2.501	0.821	1.377	0.777
Slope $\hat{\beta}_2$	1.279	4.026	1.130	1.97	0.953	0.67

Table 4.8: Random Effect Model by Liquidity. – Intercept and Slope

	low liquidity		mid liquidity		high liquidity	
Random Part	StdDev	#Firms	StdDev	#Firms	StdDev	#Firms
Intercept η_i	1.289	811	2.222	804	2.076	796
Slope θ_i^1	1.164		0.910		1.053	
Slope θ_i^2	1.362		0.215		0.268	
Fixed Part	Value	t-value-1	Value	t-value-1	Value	t-value-1
Intercept α	8.238		4.379		3.966	
Slope β_1	1.462	0.9096	2.481	2.31	1.167	0.278
Slope β_2	1.362	4.701	1.112	1.604	0.938	0.713

first 15, 30, 60 minute returns based on the same dataset as that of the first 5 minute return, and the results are summarized in Table 4.9 and Table 4.10.

When we allow both the intercept and slopes to be random, the slope coefficient of the estimated volatility during halt is greater than one in all three cases of different time horizon, indicating the price limit has a significant leverage effect on the volatility. For the slope coefficient of the estimated overnight volatility, it is not different from one in all three cases. Therefore, the mechanical effect still exists, but is not enlarged. Moreover, it is noted that the intercept is always greater than zero, and hence we have an overwhelmingly rejection to the null hypothesis. We conclude that the spillover effect varies depending on the time horizon we look at, and the effect is significant with no cooling off effect at any level to be identified. In summary, we examine the spillover effect of price limits on the opening returns (over various

Table 4.9: Random Effect Model by Time Interval. – Intercept

	15 min		30 min		60 min	
Random Part	StdDev		StdDev		StdDev	
Intercept η_i	4.559		4.485		4.862	
Residual ϵ_{ij}	12.879		12.869		15.581	
Fixed Part	Value	t-value-1	Value	t-value-1	Value	t-value-1
Intercept α	8.823		8.268		10.979	
Slope $\hat{\beta}_1$	1.041	0.26	0.862	1.150	1.129	1.142
Slope $\hat{\beta}_2$	1.113	2.63	1.110	2.62	1.038	0.745

Table 4.10: Random Effect Model by Time Interval. – Intercept and Slope

	15 min		30 min		60 min	
Random Part	StdDev		StdDev		StdDev	
Intercept η_i	3.898		3.421		5.881	
Slope θ_i^1	1.223		0.920		1.009	
Slope θ_i^2	0.256		0.252		0.305	
Fixed Part	Value	t-value-1	Value	t-value-1	Value	t-value-1
Intercept α	7.436		6.784		9.541	
Slope β_1	1.355	1.571	1.105	0.618	1.288	1.756
Slope β_2	1.157	3.204	1.155	3.163	1.095	1.890

time spans) of the next day. We find significant evidence for this effect as well as the mechanical effect, after controlling for a company's specific overnight characteristics. Contrary to what might be expected after the emergence of a consensus price, the period immediately after a trading cessation is characterized by higher levels of volatility.

Chapter 5

Conclusion

We use 5-minute intra-day price data from the Tai Wan stock exchange(TSE), where a daily price limit of 7% is imposed for all traded stocks, to study the impact of price limits on volatility. Our null hypothesis of interest is that knowing whether or not the price reaching the limits and staying there till the end of the day has no spillover effect on the return volatility of the next day. The empirical test of this hypothesis is implemented by cross sectional random effect regression model, and the expected volatility during the halt is obtained by GARCH forecasting and simulation.

If the null hypothesis is true, slopes β_1 for the estimated overnight volatility, and β_2 for the estimated volatility during a halt in the cross sectional regression should be one. If on the contrary the price limits elevate return volatility on the following day, slopes should be greater than one. The third case where βs are smaller than one then indicates cooling off effect.

The cross sectional random effect model fitted by company has a slope β_2 of the estimated volatility during a halt significantly larger than one, in either case where only the intercept is allowed to be random, or both the intercept and slopes are allowed to be random. The cross sectional random effect model fitted by sector delivers the same result. We further test the null hypothesis within each sector and find that out of 17 sectors, 15 have β_2 larger than one.

In the next step, we classify the companies into three different liquidity levels and ensure the number of companies in each category about the same. We find the test results are slightly varied according to liquidity level, with the highest liquidity companies corresponding to β values not different from one, and both mid and low liquidity companies corresponding to β values larger than one. There is no cooling off effect identified. In addition, we test the null hypothesis using the first 15-minute, 30-minute and 60-minute opening returns following trading cessations. The results reveal that as the time window examined being extended, the magnitude of the spillover effect decays slightly. Overall, the slope coefficient (β_2) of the estimated volatility during a halt is significantly greater than one; the slope coefficient (β_1) of the estimated overnight volatility is not different from one; and the intercept is greaer than zero. Again, there is no cooling off effect at any level identified. In the end, we conclude that the daily price limits in TSE are ineffective in preventing over-reaction. Price limits either mechanically increase the volatility of the opening return in a post-event day, or, they additionally have spillover effect that elevates the volatility to a even higher level. We recommend the policy makers to remove the price limits on the Tai Wan Stock Exchange.

Reference

cha:App

Andersen, T.G., T. Bollerslev, 1997, Intraday Periodicity and Volatility Persistence in Financial Markets, *Journal of Empirical Finance 4, 115-158*

Arak, M., R.E. Cook, 1997, Do Daily Price Limits Act as Magnets? The Case of Treasury Bond Futures, *Journal of Financial Services Research 12:1 5-20*

Bertero, E., C. Mayer, 1990, Structure and Performance: Global Interdependence of Stock Markets Around the Carsh of OCtober 1987, *European Economic Review, 34, 1155-1180*

Bollerslev, T.,1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics, 31, 307-327.*

Cho, D., J. Russell, R. Tsay, and G. Tiao, 2002, The Magnet Effect of Price Limits: Evidence from High Frequency Data in Taiwan, *Revised*

Fama, E.F., 1989, Perspectives on October 1987, or What Did We Learn from the Crash? in Robert W. Kamphuis, Jr. Roger C. Kormendi, and J.W.. Henry Watson, Eds.: *Black Monday and the Future of the Financial Markets* (Irwin, Homewood, Ill.)

George, T.J., C.Y. Hwang, 1995, Transitory Price Changes and Price-Limit Rules: Evidence from the Tokyo Stock Exchange, *The Journal of Financial and Quantitative Analysis, 30:2*

Greenwald, B.C., J.C. Stein, 1988, Transactional Risk, Market Crashes, and the Role of Circuit Breakers, *The Journal of Business, 64:4, 443-462*

Kim, K.A., S.G. Rhee, 1997, Price Limit Performance: Evidence from Tokyo Stock Exchange, *Journal of Finance 52, 885-901.*

Kuhn, B.A., G.J. Jurserk, and P. Locke, 1991, Do Circuit Breakers Moderate Volatility? Evidence from October 1989, *The Review of Futures Markets 10, 136-175.*

Lauterbach, Beni, U.B. Zion, 1993, Stock Market Crashes and the Performance of Circuit Breakers: Empirical Evidence, *The Journal of FInance, 48:5*

Lee, C.M.C., M.J. Ready, and P.J. Seguin, 1994, Volume, Volatility, and New York Stock Exchange Trading Halts, *Journal of Finance 49, 183-214.*

Roll, R., 1988, The International Crash of OCtober 1987, *Financial Analyst Journal, 19-35*

Subrahmanyam, A., 1994, Circuit Breakers and Market Volatility: A Theoretical Perspective, *Journal of Finance 49, 527-543*