

University of Wollongong

Research Online

Faculty of Engineering and Information
Sciences - Papers: Part A

Faculty of Engineering and Information
Sciences

1-1-2016

Scaling by the square-root-of-time rule: An empirical investigation using five market indexes

James Cameron
University of Wollongong

Chandra M. Gulati
University of Wollongong, cmg@uow.edu.au

Yan-Xia Lin
University of Wollongong, yanxia@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/eispapers>



Part of the [Engineering Commons](#), and the [Science and Technology Studies Commons](#)

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Scaling by the square-root-of-time rule: An empirical investigation using five market indexes

Abstract

The value of the theoretical constant used to scale daily volatility to both a five-day volatility estimate and a ten-day volatility estimate is compared with empirical estimates using a volatility modeling framework. Five composite stock indexes are analyzed to determine the different behaviors of scaling across markets. Both developed and emerging markets are considered, to provide additional detail to the comparisons. The results provided are considered in a value-at-risk application. While using the square-root-of-time rule on a weekly or ten-day basis is appropriate in certain cases, for time series with a linear dependence component the rule can drastically err from observed volatility levels. It is demonstrated that there are potential hazards when using the square-root-of-time rule for risk or compliance purposes.

Disciplines

Engineering | Science and Technology Studies

Publication Details

Cameron, J., Gulati, C. & Lin, Y. (2016). Scaling by the square-root-of-time rule: An empirical investigation using five market indexes. *Journal of Risk*, 19 (2), 61-80.

Research Paper

Scaling by the square-root-of-time rule: an empirical investigation using five market indexes

James Cameron, Chandra Gulati and Yan-Xia Lin

National Institute for Applied Statistics Research Australia, School of Mathematics and Applied Statistics, University of Wollongong, Northfields Avenue, Wollongong, NSW 2522, Australia;
emails: jamescam92@gmail.com, cmg@uow.edu.au, yanxia@uow.edu.au

(Received September 17, 2015; revised March 16, 2016; accepted March 24, 2016)

ABSTRACT

The value of the theoretical constant used to scale daily volatility to both a five-day volatility estimate and a ten-day volatility estimate is compared with empirical estimates using a volatility modeling framework. Five composite stock indexes are analyzed to determine the different behaviors of scaling across markets. Both developed and emerging markets are considered, to provide additional detail to the comparisons. The results provided are considered in a value-at-risk application. While using the square-root-of-time rule on a weekly or ten-day basis is appropriate in certain cases, for time series with a linear dependence component the rule can drastically err from observed volatility levels. It is demonstrated that there are potential hazards when using the square-root-of-time rule for risk or compliance purposes.

Keywords: scaling; volatility; square-root-of-time rule; GARCH model; EGARCH model.

1 INTRODUCTION

In modern finance, it is commonplace to estimate the volatility of the return on an asset over a t -day period,¹ given that the volatility of that asset's return for a single day is known.² It is simple to obtain the one-day return volatility of an asset. To then make a theoretically accurate prediction about the volatility over the t -day horizon might seem very appealing. One such approach is to apply a scaling constant to the one-day return volatility; this is known as the square-root-of-time rule:

$$\hat{\sigma}_t = \sqrt{t}\sigma, \quad (1.1)$$

where σ is the volatility of the daily return. The degree of accuracy of volatility estimates derived under various approaches relative to observed volatilities is particularly important, since volatility estimates are often a crucial input into asset and derivative pricing models and risk estimation models. There are a few key underlying assumptions for the derivation of the square-root-of-time rule. It is assumed that the changes in the logarithm of an asset price are independent and identically distributed (iid), which can be represented as

$$\log(p_n) - \log(p_{n-1}) = \varepsilon_n. \quad (1.2)$$

Here, we have a standard deviation of σ , and $\{\varepsilon_n\}$ are iid with zero mean and finite variance; it is often also assumed that these terms are normally distributed. The difference in the logarithms of an asset price after one day, $\log(p_n) - \log(p_{n-1})$, is called the one-day log return. Throughout this paper, the term "return" is used to mean "log return", and it is based on the closing price.

Similarly, the t -day return can be calculated as

$$\log(p_n) - \log(p_{n-t}) = \sum_{i=0}^{t-1} \varepsilon_{n-i}.$$

As the ε terms are independent, the variance of the t -day return is simply the sum of the t -day variances, $t\sigma^2$. Taking the square root to obtain the standard deviation (volatility), it is obvious that this is equal to $\sqrt{t}\sigma$. Further explanation and a discussion, focusing on Brownian motion, can be found in Parlar (2004).

While it is common to simplify real-world mechanisms in order to attempt to model observed processes, the square-root-of-time rule is arguably an oversimplification, with its assumptions meriting criticism. Research has shown that the iid-normal assumption regarding asset returns is often violated (Groenewold and Fraser

¹ In this case, the volatility of an asset's return is synonymous with the standard deviation of the return. It is often associated with how risky the asset is generally viewed to be.

² The rule does not have to be used on a daily basis alone. It could be applied to any time step.

2001, 2002). Doubt surrounds the ability of the square-root-of-time rule to accurately estimate volatility levels over various time frames. Namely, Diebold *et al* (1998) criticized the rule for its behavior as the time frame increases: it is easy to see that as t gets large, so too will $\hat{\sigma}_t$. Evidently this is not the case, as asset return volatility is generally bounded (Cont 2001). Conversely, Kumar (2006) found that in the case of the Indian stock exchange, scaling by the square root of time underestimates volatilities for periods of up to 120 days. Brummelhuis and Kaufmann (2007) investigated scaling of a short-term value-at-risk (VaR) to get a long-term VaR by using the square-root-of-time rule and found that, although not perfect, the $\sqrt{10}$ rule performed reasonably well. Daniélsson and Zigrand (2006) considered the performance of the square-root-of-time rule in VaR calculations for extreme financial events, also known as tail events. One of the main criticisms of VaR is its inability to estimate potential losses when large financial shocks occur. While not ruling out the use of the square-root-of-time rule overall, Daniélsson and Zigrand concluded that it was an inappropriate tool under the circumstances of their experiment. Wang *et al* (2011) also examined the performance of the square-root-of-time rule for bias in VaR.

Research conducted on the square-root-of-time rule has predominately focused on the theoretical validity of the practice. This question understandably remains in the spotlight. To date, it seems little work has been dedicated to exploring the scaling coefficient empirically. The empirical nature of the rule deserves investigation, as this will help illuminate its observed behavior. The primary goal of this paper is to determine the empirical values of the scaling coefficient over a five-day period and a ten-day period across five different capitalization stock indexes. This is achieved, for the most part, through a volatility modeling framework. Daily log returns are calculated for the five indexes over the period 2003–13, and the return volatility is modeled using two different generalized autoregressive conditional heteroscedasticity (GARCH) models, namely the original GARCH model of Bollerslev (1986) and the exponential GARCH (EGARCH) model of Nelson (1991). Regression analysis is the primary method of calculating the scaling coefficients. A secondary method is derived to provide both an alternate approach and a comparison between values. Using five major stock indexes from different markets represents a global approach and will allow discrepancies in market behaviors to be observed. Four of the five indexes belong to developed markets, while the fifth represents an emerging market. Emerging markets are well known for their volatile behavior relative to developed markets (Santis and Imrohoroglu 1997; Selçuk 2004). This comparison between market types will add to the emerging market literature, with a focus on volatility scaling. Additionally, we investigate several closely related questions regarding volatility scaling and test the presence of a “day of the week” effect for scaling estimates, as well as the nature of the scaling constant over different time periods.

2 METHODOLOGY

Eleven years of daily closing prices for five capitalization indexes were analyzed: the Standard and Poor's (S&P) 500, Australian Securities Exchange (ASX) 200, Nikkei 225, Financial Times Stock Exchange (FTSE) 100 and BM&F Bovespa. The data analyzed is from the period 2003–13, which yielded 2893 daily observations, 573 weekly observations and 286 ten-day observations. The selected time frame captures the turmoil experienced during the financial crisis. It is certainly of interest to determine the effect that this significant market event had on the nature of volatility scaling in this period. Where there are missing values in each time series due to closures, the previous day's closing price has simply been used as the missing observation. This was done to provide continuity of the data and to permit consistent analysis when comparing daily volatility levels with their weekly counterparts.³

Most of the indexes analyzed belong to developed markets, while the BM&F Bovespa from the São Paulo exchange represents the only emerging market in the data set. The BM&F was specifically included to add an extra dimension to the analysis, as it is generally viewed that emerging markets can offer high returns to investors, but these are often associated with high levels of risk (Šević 2005). It is entirely possible that the more volatile emerging market may have a different scaling climate than the developed markets. A summary of descriptive statistics for the various return series can be seen in Table 1.

Two methods for estimating the scaling coefficient are proposed in this paper. To the authors' knowledge, neither has been employed in previous research. The methods proposed make two different assumptions about the type of volatility used in each approach. One method assumes that the variance of the asset return is constant over time. The idea of static return volatility does conflict with known volatility characteristics. However, as only five-day and ten-day periods are analyzed, this open approach should not be overly erroneous in its estimation. This method is arguably the simplest of the two, in both its assumptions and implementation. The second method employs volatility models from the GARCH family, thus treating the volatility as a conditional volatility.

2.1 Method 1

The first proposed method utilizes the autocorrelation values of the daily return time series to estimate the scaling coefficient. A key assumption of this method is that the

³ Different approaches in assigning a value to the missing value might result in different impacts on volatility estimation. Deciding which approach to adopt is an interesting open question.

TABLE 1 Descriptive statistics for daily, weekly and ten-day log return series for 2003–13.

(a) Daily returns					
	ASX 200	S&P 500	Nikkei 225	FTSE 100	BM&F Bovespa
Obs.	2867	2867	2867	2867	2867
Mean	0.000196	0.000248	0.000220	0.000182	0.000520
Median	0.000223	0.000476	0.000000	0.000133	0.000172
Maximum	0.056282	0.109572	0.132346	0.093842	0.136766
Minimum	-0.087043	-0.094695	-0.121110	-0.092646	-0.120961
Variance	0.000114	0.000156	0.000228	0.000142	0.000309
Skewness	-0.466162***	-0.311863***	-0.595917***	-0.119784***	-0.083193**
Kurtosis	5.945726***	11.195644***	8.645413***	8.464411***	5.679719***
(b) Weekly returns					
	ASX 200	S&P 500	Nikkei 225	FTSE 100	BM&F Bovespa
Obs.	573	573	573	573	573
Mean	0.000968	0.001194	0.001092	0.000908	0.002540
Median	0.002973	0.002237	0.003639	0.002803	0.005136
Maximum	0.102771	0.129509	0.180167	0.142649	0.261933
Minimum	-0.120775	-0.149083	-0.235397	-0.130291	-0.292621
Variance	0.000550	0.000684	0.001219	0.000651	0.001679
Skewness	-0.643579***	-0.349591***	-0.970328***	-0.304817***	-0.562234***
Kurtosis	3.607046***	5.504757***	7.329576***	4.211273***	7.647691***
(c) Ten-day returns					
	ASX 200	S&P 500	Nikkei 225	FTSE 100	BM&F Bovespa
Obs.	286	286	286	286	286
Mean	0.001898	0.002367	0.002097	0.001791	0.005078
Median	0.005622	0.006474	0.006460	0.004664	0.009223
Maximum	0.118202	0.195882	0.147901	0.109620	0.144174
Minimum	-0.144642	-0.177903	-0.150964	-0.156088	-0.208839
Variance	0.001000	0.001228	0.001868	0.001068	0.002707
Skewness	-0.731581***	-0.696933***	-0.465726***	-0.795471***	-0.583923***
Kurtosis	3.318299***	7.188231***	1.062386***	3.021099***	1.482631***

*, ** and *** denote statistical significance at the 10%, 5% and 1% confidence levels, respectively.

series is weakly stationary.⁴ The variance of the weekly returns can be equated to the variance of the sum of one-day returns, where

$$\text{var}(r_{w_j}) = \text{var}(r_{j,1} + r_{j,2} + r_{j,3} + r_{j,4} + r_{j,5}).$$

Here, r_{w_j} represents the weekly return for week j , and $r_{j,i}$ represents the return on day i in week j . The l -step covariance can be estimated as the product of the lag- l autocorrelation coefficient and the one-day variance, ie, $\text{cov}(r_{t+l}, r_t) = \rho_l \sigma^2$,

$$\text{var}(w) = (5 + 8\rho_1 + 6\rho_2 + 4\rho_3 + 2\rho_4)\sigma^2,$$

and the scaling constant can be estimated by

$$c_5^* = \sqrt{5 + 8\rho_1 + 6\rho_2 + 4\rho_3 + 2\rho_4}. \quad (2.1)$$

The ρ values are estimated as the lag- l autocorrelations using the data. Formula 2.1 can be simply extended to calculate the coefficient over a ten-day period:

$$c_{10}^* = \sqrt{10 + 18\rho_1 + 16\rho_2 + \dots + 4\rho_8 + 2\rho_9}. \quad (2.2)$$

If r_t follows an autoregressive AR(1) model

$$r_t = \mu + \phi r_{t-1} + a_t,$$

where a_t is a white noise series with zero mean and variance σ_a^2 , then ρ_l can be calculated as

$$\rho_l = \phi^l. \quad (2.3)$$

This can then be applied to provide an alternative to the c^* formula:

$$c_n^{**} = \sqrt{n + 2(n-1)\phi + \dots + 2\phi^{n-1}}. \quad (2.4)$$

By fitting an AR(1) model to the data, the estimated value of ϕ can be used in (2.4).

Alternatively, if r_t were to follow a moving average MA(1) model

$$r_t = \mu - a_t - \theta a_{t-1},$$

then $\rho_1 = -\theta/(1 + \theta^2)$, but ρ_k for $k > 1$ would simply equal zero, thus greatly reducing the number of terms in the c^{**} equation:

$$c_n^{**} = \sqrt{n + 2(n-1)\frac{-\theta}{1 + \theta^2}}. \quad (2.5)$$

The most appealing feature of this approach is its overall simplicity. The results for this method are presented in Tables 3 and 4. They show that over a five-day period this method produces estimates similar to those estimated by the approach presented in Section 2.2, but, when extended to a ten-day window, the similarity between the two estimates fades.

⁴ This includes the assumptions of the variance of returns is equal for each day and that the individual l -step covariances only depend on l .

TABLE 2 Summary of scaling coefficients using the regression method over a five-day period.

Index	C_{GARCH}	C_{EGARCH}
ASX 200	2.15491***	2.16774***
R^2	0.9493	0.9654
RSE	0.0054	0.004682
S&P 500	1.94350***	1.97792***
R^2	0.9255	0.9548
RSE	0.007309	0.005333
Nikkei 225	2.18523***	2.22115***
R^2	0.9486	0.9548
RSE	0.0077	0.007159
FTSE 100	2.05956***	2.10547***
R^2	0.9106	0.9606
RSE	0.006888	0.004941
BM&F Bovespa	2.27029***	2.27641***
R^2	0.9573	0.9711
RSE	0.008312	0.006737

*** Statistical significance at 1% confidence level. RSE denotes residual standard error, used here as a measure of the variability (standard deviation) of the predicted parameters.

TABLE 3 Scaling coefficients estimated using the autocorrelation of the daily returns, c^* , compared with the results given by setting $\rho_k = \phi^k$, c^{**} , and the regression coefficients given by EGARCH volatility estimates.

	ASX 200	S&P 500	Nikkei 225	FTSE 100	BM&F Bovespa
c_5^*	2.133647	2.133608	2.135063	2.053631	2.118409
c_5^{**}	2.178937	2.034837	2.207324	2.138324	2.210066
C_{EGARCH}	2.17872	1.97792	2.22115	2.10547	2.27641

2.2 Method 2

The second proposed method employs a volatility model framework, where the volatility of the time series is modeled using GARCH and EGARCH. The method is a two-step process: the volatility of the underlying time series is modeled, and then regression analysis is used to calculate the scaling coefficient. To perform the regression the volatility estimates from the five-day and ten-day series, respectively, are regressed against the volatility estimates of the daily return series. An intercept term is not included in the regression, as we attempt to explain t -day volatility purely in terms of one-day volatility. The volatility of the return series is modeled using the

TABLE 4 Summary of scaling coefficients over a ten-day period.

Index	C_{GARCH}	C_{EGARCH}	c_{10}^{**}
ASX 200	2.81259***	2.89933***	3.08004
R^2	0.9131	0.943	
RSE	0.009487	0.007553	
S&P 500	2.63979***	3.22338***	3.06196
R^2	0.8661	0.8826	
RSE	0.01292	0.01355	
Nikkei 225	2.64685***	2.73380***	3.30987
R^2	0.8686	0.8751	
RSE	0.01568	0.01535	
FTSE 100	2.60709***	2.85206***	2.97318
R^2	0.8622	0.9306	
RSE	0.01251	0.008973	
BM&F Bovespa	2.83994***	2.90528***	3.05935
R^2	0.9056	0.9352	
RSE	0.01596	0.01317	

***Statistical significance at 1% confidence level. RSE denotes residual standard error.

GARCH(m, s) volatility model (Bollerslev 1987)⁵

$$r_t = \mu_t + a_t, \quad a_t = \sigma_t \varepsilon_t,$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^m \alpha_i a_{t-i}^2 + \sum_{j=1}^s \beta_j \sigma_{t-j}^2 \quad (2.6)$$

subject to

$$m \geq 0, \quad s \geq 0,$$

$$\alpha_0, \alpha_i \geq 0 \quad \text{for } i = 1, 2, \dots, m,$$

$$\beta_i \geq 0 \quad \text{for } i = 1, 2, \dots, s,$$

$$\sum_{i=1}^{\max(m,s)} \alpha_i + \beta_i < 1,$$

and the EGARCH(m, s) volatility model (Nelson 1991)

$$\ln(\sigma_t^2) = \alpha_0 + \sum_{i=1}^m \alpha_i \frac{|a_{t-i}| + \gamma_i a_{t-i}}{\sigma_{t-i}} + \sum_{j=1}^s \beta_j \ln(\sigma_{t-j}^2). \quad (2.7)$$

⁵ μ_t is a constant term unless an autoregressive or moving component is included in the model for r_t . It should be noted that a misspecified mean function will lead to biased and inconsistent estimates of parameters in the conditional variance.

The EGARCH model is viewed as the more robust of the two models and is thus included to determine whether any difference in the scaling coefficient is observed, the idea being that perhaps a better model will produce a more accurate estimate. For the error distribution of the models, the t -distribution has been used in nearly all cases.⁶ This was chosen over the normal distribution due to the nonnormal qualities, such as excess kurtosis and negative skewness, displayed in the data.

To calculate the scaling coefficient for a five-day window, every weekly conditional volatility estimate (for the corresponding week ahead) is regressed against every fifth daily conditional volatility estimate.⁷ This is done to calculate the coefficient over a Monday-to-Monday week:

$$\begin{aligned}\hat{\sigma}_{w,t} &= c\hat{\sigma}_{d,t} + \varepsilon_t, \\ \hat{\sigma}_w &= E(\hat{\sigma}_{w,t}) = E(c\hat{\sigma}_{d,t} + \varepsilon_t) = c\hat{\sigma}_d,\end{aligned}$$


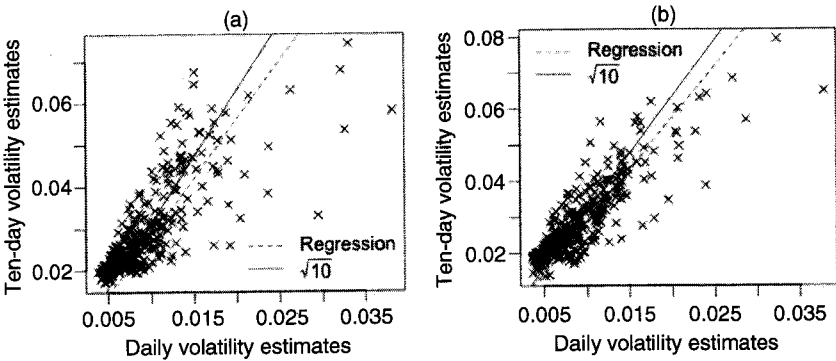
where $\hat{\sigma}_w$ and $\hat{\sigma}_d$ represent the mean of weekly and daily volatility estimates, respectively. Naturally, this is extended to a ten-day window for $t = 10$. If scaling by the square root of time were to hold true, it would be expected that the value of c would be approximately equal to $\sqrt{5} \approx 2.236$ for the five-day case and $\sqrt{10} \approx 3.162$ in the ten-day case. In cases where the volatility of a return series is modeled using an integrated GARCH (IGARCH) volatility model, we use a similar method to the regression approach proposed herein for estimating volatility levels over a t -day period. The volatility forecasts provided by the IGARCH model form a straight line with a constant slope (for more information, see Tsay (2010)). This approach is appealing due to its straightforward implementation and, notably, the ability to display the various scaled volatility results through a regression plot. The results can be seen in Tables 2 and 4, respectively.

3 RESULTS

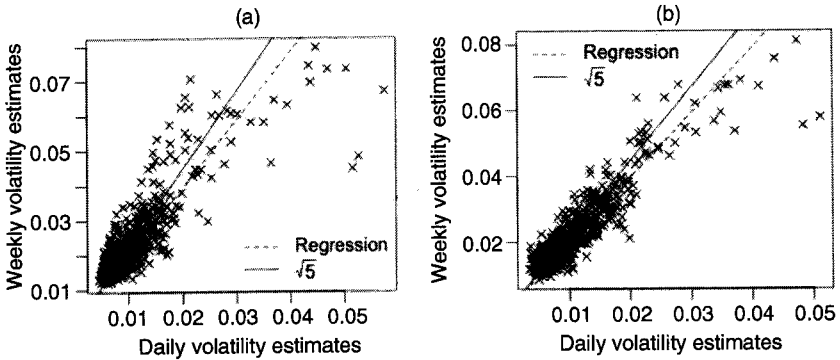
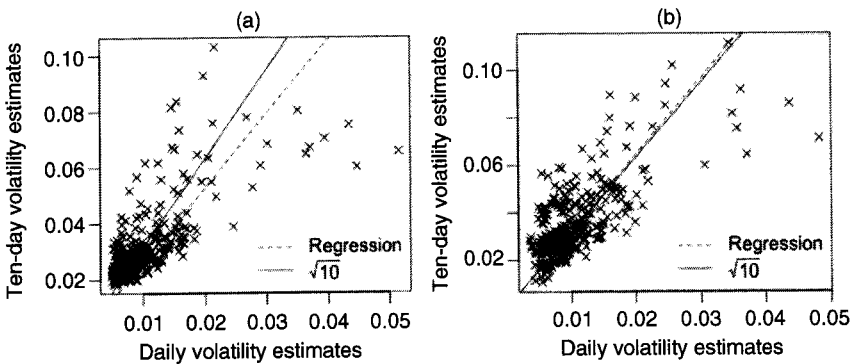
A considerable range of values is observed for the scaling estimates as calculated by the regression method in Table 2. Overall, the S&P 500 index gives the lowest estimates for both models, while the BM&F produces the highest overall estimates in both cases. The S&P 500 and FTSE 100 indexes scored the two lowest scaling coefficients, which is interesting given that both series required a first-order AR term or MA component in the return equations for both the daily and weekly GARCH models and daily EGARCH models: a notable contrast from the other indexes. The signs of the coefficients of these additional AR or MA terms are negative; Tsay (2010)

⁶ In a few minor instances, a generalized error distribution (GED) was used when the shape parameter in the model output produced an insignificant coefficient.

⁷ The estimates provided by the volatility models are conditional volatility estimates.


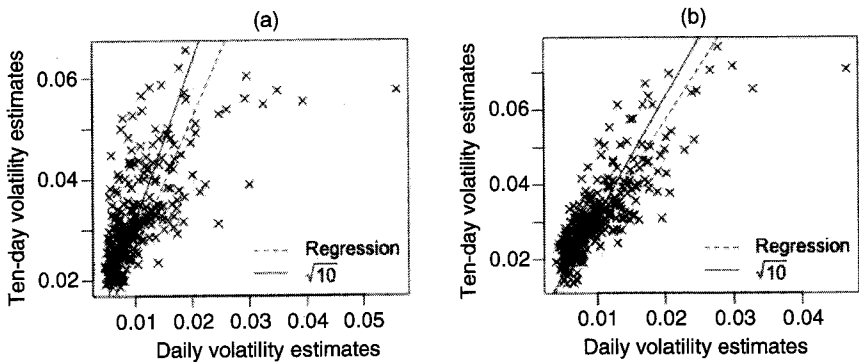
FIGURE 1 Five-day ASX (a) GARCH regression and (b) EGARCH regression.**FIGURE 2** Ten-day ASX (a) GARCH regression and (b) EGARCH regression.

mentions that this is not an uncommon property of financial time series. Simulations (not provided here) were run to determine the effect of both a positive AR presence and a negative AR presence in the return series. The inclusion of a positive lag -1 , or AR(1), coefficient (denoted by ϕ) had the effect of raising the scaling coefficient estimates above $\sqrt{5} \approx 2.236$. As the ϕ increased, so too did the scaling coefficient estimates as calculated by the two abovementioned methods. A negative ϕ coefficient had the opposite effect. Its presence lowered estimates below $\sqrt{5}$ and, as the coefficient decreased in value, so too did the scaling estimates. From (2.4), it can be seen that if $\phi = 0$ and $n = 5$, the scaling coefficient should be exactly equal to $\sqrt{5}$. It is obvious that in all cases the regression coefficients using volatility estimates from EGARCH

FIGURE 3 Five-day S&P 500 (a) GARCH regression and (b) EGARCH regression.**FIGURE 4** Ten-day S&P 500 (a) GARCH regression and (b) EGARCH regression.

models are higher and produce higher R^2 values.⁸ In general, the EGARCH plots have fewer, less extreme outliers, with the coordinates generally forming a tighter fit around the regression line. Graphs of the regression plots are shown in Figures 1–10. (Details of the GARCH and EGARCH fitted models are provided in the online appendix.) This evidence suggests that the EGARCH results are the better and more accurate of the two models.

⁸The R^2 value is a measure of how closely the fitted equation models the empirical data. In this instance, a higher R^2 is desirable, as this means that more of the variability in the data is explained. The R^2 value can be seen as an indication of the model's predictive ability.

FIGURE 5 Five-day FTSE 100 (a) GARCH regression and (b) EGARCH regression.**FIGURE 6** Ten-day FTSE 100 (a) GARCH regression and (b) EGARCH regression.

Regarding the regression plots in Figures 1–10, most of the largest outliers lie well beneath the $\sqrt{5}$ line. As the weekly volatility estimates increase, it appears we are less likely to observe plot points above the $\sqrt{5}$ line. When daily volatility estimates lie within the 0.02–0.03 range, weekly estimates are very often observed above their theoretical value. But as the daily estimates increase, this virtually ceases, almost as if the behavior of weekly estimates changes once daily volatility estimates hit 0.03. The most severe outliers on the regression plots generally occur beneath the regression line. Removal of these outliers naturally results in an increase in the gradient of the regression line. In certain cases, the removal of the three most significant points alone resulted in a scaling coefficient estimate almost directly equal to the theoretical value.

FIGURE 7 Five-day Nikkei 225 (a) GARCH regression and (b) EGARCH regression.

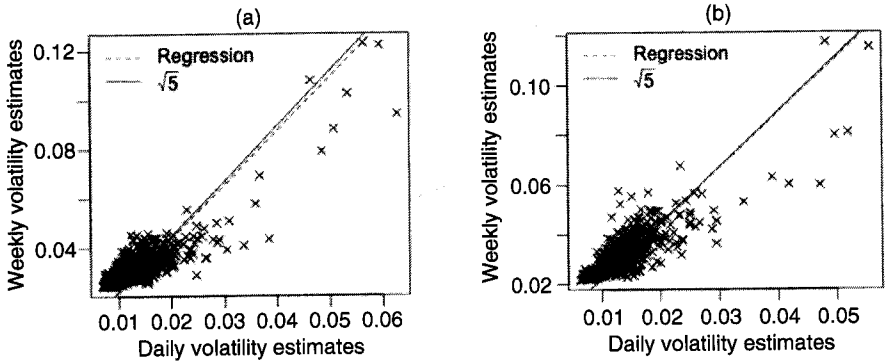
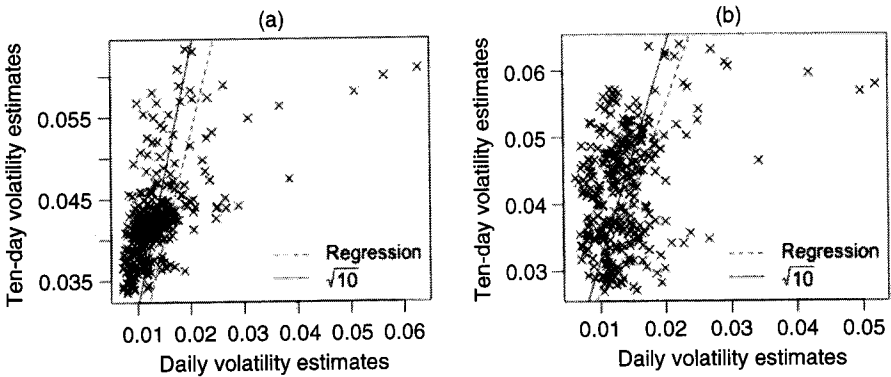

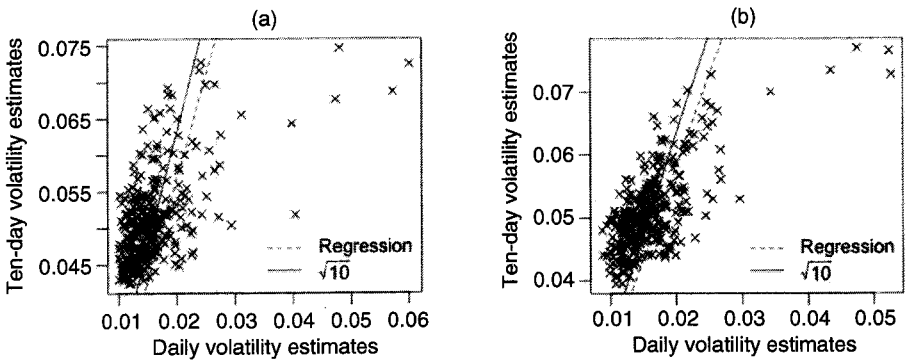


FIGURE 8 Ten-day Nikkei 225 (a) GARCH regression and (b) EGARCH regression.



The results obtained here are varied. In certain cases, it seems that scaling a daily volatility value by $\sqrt{5}$ would be perfectly appropriate, yet in other instances such a practice would overestimate weekly volatility levels. This result is out of line with the expectation of Diebold *et al* (1998); however, it is not entirely unsurprising, as a five-day period is not exactly the “large” time period we warned about earlier. Certain results seen here seem to suggest that, in particular cases, scaling by $\sqrt{5}$ is an appropriate and reasonably accurate practice. Results to the contrary occur when a series requires an AR or MA component in the return equation. The observed coefficients for the AR and MA terms were negative, which seemingly corresponds

FIGURE 9 Five-day BM&F Bovespa (a) GARCH regression and (b) EGARCH regression.**FIGURE 10** Ten-day BM&F Bovespa (a) GARCH regression and (b) EGARCH regression.

to a decrease in the value of the scaling coefficient estimates. This may suggest a relationship between the linear dependence of a series and its volatility.

The results given by (2.1) can be seen in Table 3. The variation witnessed in the regression coefficients is seemingly absent in the results given by c^* , but it can be observed to some extent in the c^{**} values. Admittedly, the overall consistency of c^* is surprising. Does this result suggest that the true empirical scaling coefficient lies somewhere around 2.1? This may be unlikely, as setting $\rho_l = \phi^l$ seemingly permits (2.1) to produce results that mimic those of the EGARCH regression more closely. This alternative method of calculating the scaling coefficient is appealing in its simplicity and ease of implementation. From these values, the actual autocorrelations

TABLE 5 Mean of daily variance estimates as measured by an EGARCH model.

Day	ASX 200	S&P 500	Nikkei 225	FTSE 100	BM&F Bovespa
Monday	0.0094201	0.0105803	0.0137075	0.0102565	0.0161775
Tuesday	0.0094185	0.0105684	0.0137171	0.0102595	0.0161800
Wednesday	0.0094030	0.0103754	0.0136769	0.0102297	0.0161738
Thursday	0.0093990	0.0105095	0.0136366	0.0102927	0.0162065
Friday	0.0094472	0.0104524	0.0136744	0.0103208	0.0162547
Mean	0.0094175	0.0104972	0.0136825	0.0102718	0.0161985
Range	0.0000482	0.0002049	0.0000805	0.0000911	0.0000809

Source: Berument and Kiyamaz (2001).

do not appear to produce desirable results over a five-day period. The method adopted in the c^{**} calculations seems to be the most preferable approach at this stage.

Further deviations from the theoretical \sqrt{t} value are observed in Table 4. In all cases but one, the empirical coefficients fall beneath $\sqrt{10}$. This certainly clashes with the expectation of overestimation as the time horizon is increased. Unfortunately, the Brazilian ten-day series did not display ARCH effects; thus, GARCH models could not be used to model the volatility of the series. The R^2 values dropped across the board due to the larger variation in volatility estimates. This certainly provides evidence to support our argument that the accuracy of the square-root-of-time rule diminishes as the time horizon increases. Similarly to the five-day scenario, the main outliers on the regression plots correspond to high daily volatility estimates with unexpectedly low ten-day volatility values. The c^{**} values do not mimic the EGARCH regression coefficients as closely as in the five-day case: the difference between the two types of estimates has grown. Interestingly, the inclusion of an autoregressive moving average component in the daily return series does not appear to have the same effect as in the five-day case. This may be because the effect of these variables diminishes over time.

3.1 Other results

Two additional issues relating to volatility scaling, not included here, were investigated. A test-case rolling window analysis was performed on the S&P 500 index to determine whether or not the indexes' scaling coefficient changed over time. In addition, the question of whether the scaling coefficient is dependent upon the start and end days of the time period being investigated was briefly explored: this is not unlike the "day of the week" effects observed in asset returns (French 1980; Keim and Stambaugh 1984). To perform the rolling window analysis, an approximate six-year window was constructed, rolling forward approximately one year with each iteration,

thus giving six different windows. Only a five-day period was considered at this stage. While this analysis was only performed on a single index, the results certainly suggest that the scaling coefficient is dynamic. Testing the influence of the days on which a five-day window started and finished also used only the S&P 500 data. This was because the difference between the highest daily average volatility and lowest average daily volatility was greatest for this data set. Average daily volatility values were taken from an EGARCH model, and the scaling coefficients were calculated for the days with the highest and lowest average volatilities, respectively, over a five-day period. For the S&P 500, Monday had the highest average volatility, and Wednesday the lowest. The average daily volatilities for each index can be seen in Table 5. Some variation was observed between the coefficients calculated over the two different five-day windows. The coefficient for the Monday–Monday week was greater than that of the Wednesday–Wednesday week. While further research is definitely required, a relationship between the volatility on a given day and the scaling coefficient over a time frame starting and ending on that day may exist. Namely, the lower the average volatility, the smaller the scaling coefficient estimate. As is observed in Table 5, Berument and Kiyamaz (2001) also found that there was no consensus on an international level regarding which particular days had the highest and lowest volatility.

3.2 Implications

One important application of the square-root-of-time rule can be observed in VaR calculations. The Basel II and Basel III accords both recommend a ten-day holding period for certain practices, in which the one-day VaR is multiplied by $\sqrt{10}$. While there have been several criticisms of the overuse of VaR in practice (Chen 2014), and a subsequent move toward using expected shortfall in certain instances, some concern still remains regarding the use of the \sqrt{t} component in the VaR equation, as this constant directly affects the VaR estimate

$$\text{VaR} = -z_{\alpha}\sigma v\sqrt{t}, \quad (3.1)$$

where v , σ and z_{α} represent the US dollar value of the portfolio, the standard deviation of the portfolio returns and the upper 100α th quantile of the standard normal distribution. In 1996, the Basel Committee on Banking Supervision stated that one-day volatility values can be scaled by $\sqrt{10}$, or a volatility estimate can be provided after modeling the volatility of the ten-day returns for the portfolio in question (see Daniéls-son and Zigrand 2006). One particular problem may present itself when attempting to model the volatility of n -day returns, in that the series may not have an adequate number of observations to produce a significant volatility model. From the results provided here, modeling the volatility of the ten-day return series would be the preferred method of the two. However, neither practice is seemingly without flaws.

For the ten-day returns, the coefficients estimated by the regression fall short of the theoretical value of $\sqrt{10}$ in nearly every case. This may at first appear to suggest that scaling by the square root of time is reasonable, as $\sqrt{10}$ is larger than the estimated coefficients. However, as VaR is a risk measure, we would not want to underestimate a potential loss of capital. The regression plots display many points that lie above both the regression and $\sqrt{10}$ lines, which would correspond to an underestimation in predicted volatility levels, the effect of which could be very severe. For the ASX data, a slope of magnitude $\sqrt{30}$ was required so that all points lay underneath the regression line. In the S&P case, a magnitude of around $\sqrt{80}$ was required. This is confirmation of the dangers of using $\sqrt{10}$, as coefficients two to three times larger would be more appropriate. The practice of multiplying VaR estimates by 2 or 3 is mentioned in Chen (2014). Our results suggest that this practice is not a cautionary measure, but perhaps a necessary one. Of course, VaR estimates could be multiplied by some constant on a case-by-case basis, as indicated by the difference between the slopes required in the ASX and S&P cases. While alternative scaling methods have been proposed by Gençay *et al* (2001), the shortfalls of the square-root-of-time rule could perhaps be overcome by doubling or tripling the constant's value. This, however, could be seen as a rule of thumb for a rule of thumb. More concrete approaches may be required for regulatory purposes.

4 CONCLUSION

It has been suggested in the literature that scaling by the square root of time both over- and underestimates volatility levels for various time horizons. Scaling by $\sqrt{5}$ would overestimate the weekly volatility level for three out of our five indexes for the entire eleven-year period. For the remaining two indexes, scaling by $\sqrt{5}$ does not appear to be an erroneous practice, on average. We demonstrated that the presence of an autoregressive or moving-average component in the return equation influences the scaling coefficient estimate. For the FTSE 100 and S&P 500, which both required a negative AR(1) or MA(1) term in their daily return equations, the regression coefficients were significantly smaller than those of their counterparts. Conversely, this phenomenon was not observed in the results for the ten-day series. As the time horizon increased, so too did the variation in regression coefficients. The results indicate that, on average, $\sqrt{10}$ is larger than the estimated coefficient, as suggested by Kumar (2006). It appears that, as the time horizon increases, the square-root-of-time rule becomes less accurate and ultimately less relevant. In the case of large horizons, other methods developed by Selçuk (2004), Gençay *et al* (2001) and de Mattos Neto *et al* (2011) might prove to be more fitting.

We tested a second potential method for calculating the scaling coefficient using the return series autocorrelation values. The results of this alternative approach for

the five-day returns are mixed. For the ten-day series, the overall accuracy certainly appears to have eroded slightly. It is known that return volatility is not constant over time. Rolling window analysis was performed on the S&P 500 index to deem if and how scaling coefficients change over time. In line with expectation, the scaling coefficients calculated for the six time windows did vary considerably. Under the rolling window approach, the scaling coefficients could be considered as a path of realization of a time series. Forecasting the scaling coefficient for next rolling window should be of interest. Further analysis was conducted that looked into the presence of a “weekday effect” influencing scaling estimates. While we explored just a single case, the result was certainly indicative of such an effect. This result may hint at a relationship between days with lower average volatilities, which correspond to lower volatilities for weekly periods starting on that weekday. Further cases should be studied looking at both different days and different indexes to obtain a better understanding of how prevalent this potential effect is. This would simultaneously determine how weekly volatility levels differ when different starting days are used.

While criticism remains around VaR in practice, it still sees use in certain applications. If daily volatility estimates are to be scaled by $\sqrt{10}$, significant underestimations in VaR will be observed. It is recommended that estimates be doubled or possibly tripled if necessary to accommodate the occurrence of a volatility shock. Yet this somewhat primitive suggestion weakens the precision of the approach. An alternative methodology may be required.

Finally, as capitalization indexes were exclusively analyzed, the results represent the nature of scaling for large capitalization, highly liquid stocks. It is certainly of interest to see how daily volatility estimates for lower capitalization stocks scale. This would give the general idea of the effect a stock’s liquidity would have on its scaling coefficient, an issue not explicitly explored herein. Additionally, individual stocks could be analyzed, as could sector indexes. Another extension worth actively pursuing is the inclusion of different time horizons. As only five-day and ten-day windows were examined, the values and accuracy of scaling coefficients over larger time horizons are unknown. As the time horizons increase, however, more extreme deviations from the theoretical values may be observed, as described by Diebold *et al* (1998).

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

The authors thank Michael McCrae for his valuable comments. They also appreciate two reviewers’ useful and constructive comments.

REFERENCES

- Berument, H., and Kiyamaz, H. (2001). The day of the week effect on stock market volatility. *Journal of Economics and Finance* **25**(2), 181–193 (<http://doi.org/c68mv8>).
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **31**(3), 307–327 (<http://doi.org/bdw3mq>).
- Bollerslev, T. (1987). A conditional heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* **69**(3), 542–547 (<http://doi.org/bfdkrv>).
- Brummelhuis, R., and Kaufmann, R. (2007). Time-scaling of value-at-risk in GARCH(1,1) and AR(1)-GARCH(1,1) processes. *The Journal of Risk* **9**(4), 39–94 (<http://doi.org/br9t>).
- Chen, J. M. (2014). Measuring market risk under the Basel accords: VaR, stressed VaR and expected shortfall. *IEB International Journal of Finance* **8**, 184–201 (<http://bit.ly/2bW3O05>).
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* **1**, 223–236 (<http://doi.org/d4dh8f>).
- Daniélsson, J., and Zigrand, J.-P. (2006). On time-scaling of risk and the square-root-of-time rule. *Journal of Banking and Finance* **30**(10), 2701–2713 (<http://doi.org/bzfxwx>).
- de Mattos Neto, P. S. G., Silva, D. A., Ferreira, T. A. E., and Cavalcanti, G. D. C. (2011). Market volatility modeling for short time window. *Physica A* **390**, 3444–3453 (<http://doi.org/cn2f4v>).
- de Santis, G., and Imrohoroğlu, S. (1997). Stock returns and volatility in emerging financial markets. *Journal of International Money and Finance* **16**, 561–579 (<http://doi.org/cmkj3q>).
- Diebold, F., Hickman, A., Lnoue, A., and Schuermann, T. (1998). Scale models. *Risk* **11**(1), 104–107.
- French, K. R. (1980). Stock returns and the weekend effect. *Journal of Financial Economics* **8**(1), 55–69 (<http://doi.org/b4pvnz>).
- Gençay, R., Selçuk, F., and Whitcher, B. (2001). Scaling properties of foreign exchange volatility. *Physica A* **289**, 249–266 (<http://doi.org/bfk93h>).
- Groenewold, N., and Fraser, P. (2001). Tests of asset-pricing models: how important is the iid–normal assumption? *Journal of Empirical Finance* **8**(4), 491–510 (<http://doi.org/c4cpfc>).
- Groenewold, N., and Fraser, P. (2002). Violation of the iid–normal assumption: effects on tests of asset-pricing models using Australian data. *International Review of Financial Analysis* **11**(4), 491–510 (<http://doi.org/bjbb67>).
- Keim, D. B., and Stambaugh, R. F. (1984). A further investigation of the weekend effect in stock returns. *Journal of Finance* **39**(3), 819–835 (<http://doi.org/bp78>).
- Kumar, S. S. S. (2006). Square root of time rule and the Indian stock market. *South Asian Journal of Management* **13**(4), 101–107 (<http://bit.ly/2clBCY9>).
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* **59**(2), 347–370 (<http://doi.org/bz6t2m>).
- Parlar, M. (2004). A simplified treatment of Brownian motion and stochastic differential equations arising in financial mathematics. *PRIMUS: Problems, Resources, and Issues in Mathematics Undergraduate Studies* **14**(3), 269–287 (<http://doi.org/dwm4n5>).

- Selçuk, F. (2004). Financial earthquakes, aftershocks and scaling in emerging stock markets. *Physica A* **333**, 306–316 (<http://doi.org/ftb7jz>).
- Šević, Ž. (2005). Emerging markets: preferences, risks, performance.... *Economic Change and Restructuring* **38**, 1–10 (<http://doi.org/fwk6j4>).
- Tsay, R. S. (2010). *Analysis of Financial Time Series*, 3rd edn. Wiley (<http://doi.org/bxgzwt>).
- Wang, J.-N., Cheng, Y.-P., and Yeh, J.-H. (2011). How accurate is the square-root-of-time rule at scaling tail risk: a global study. *Journal of Banking and Finance* **35**(5), 1158–1169 (<http://doi.org/cw8np2>).