

Value at Risk Models in the Indian Stock Market

by

Jayanth R. Varma

Working Paper No. 99-07-05
July 1999

The main objective of the working paper series of the IIMA is to help faculty members to test out their research findings at the pre-publication stage.



INDIAN INSTITUTE OF MANAGEMENT
AHMEDABAD 380 015
INDIA

Value at Risk Models in the Indian Stock Market

Abstract

This paper provides empirical tests of different risk management models in the Value at Risk (VaR) framework in the Indian stock market. It is found that the GARCH-GED (Generalised Auto-Regressive Conditional Heteroscedasticity with Generalised Error Distribution residuals) performs exceedingly well at all common risk levels (ranging from 0.25% to 10%). The EWMA (Exponentially Weighted Moving Average) model used in J. P. Morgan's RiskMetrics® methodology does well at the 10% and 5% risk levels but breaks down at the 1% and lower risk levels. The paper then suggests a way of salvaging the EWMA model by using a larger number of standard deviations to set the VaR limit. For example, the paper suggests using 3 standard deviations for a 1% VaR while the normal distribution indicates 2.58 standard deviations and the GED indicates 2.85 standard deviations. With this modification the EWMA model is shown to work quite well. Given its greater simplicity and ease of interpretation, it may be more convenient in practice to use this model than the more accurate GARCH-GED specification. The paper also provides evidence suggesting that it may be possible to improve the performance of the VaR models by taking into account the price movements in foreign stock markets.

Value at Risk Models in the Indian Stock Market

In volatile financial markets, both market participants and market regulators need models for measuring, managing and containing risks. Market participants need risk management models to manage the risks involved in their open positions. Market regulators on the other hand must ensure the financial integrity of the stock exchanges and the clearing houses by appropriate margining and risk containment systems.

The successful use of risk management models is critically dependent upon estimates of the volatility of underlying prices. The principal difficulty is that the volatility is not constant over time - if it were, it could be estimated with very high accuracy by using a sufficiently long sample of data. Thus models of time varying volatility become very important. Practitioners and econometricians have developed a variety of different models for this purpose. Whatever intuitive or theoretical merits any such model may have, the ultimate test of its usability is how well it holds up against actual data. Empirical tests of risk management models in the Indian stock market are therefore of great importance in the context of the likely introduction of index futures trading in India.

Data

The data used in this study consists of daily values of the National Stock Exchange's NSE-50 (Nifty) index. The NSE has back-calculated this index for the period prior to the formation of the NSE by using the prices on the Bombay Stock Exchange.

Sample Period

The data period used is from July 1, 1990 to June 30, 1998. The long sample period reflects the view that risk management studies must attempt (wherever possible) to cover at least two full business cycles (which would typically cover more than two interest rate cycles and two stock market cycles). It has been strongly argued on the other hand that studies must exclude the securities scam of 1992 and must preferably confine itself to the period after the introduction of screen based trading (post 1995).

The view taken in this study is that the post 1995 period is essentially half a business cycle though it includes complete interest rate and stock market cycles. The 1995-97 period is also an aberration in many ways as during this period there was a high positive autocorrelation in the index. High positive autocorrelation violates the weak form efficiency of the market and is suggestive of an administered market; for example, it is often seen in a managed exchange rate market. The following table shows that the autocorrelation in the stock market was actually low till about mid 1992 and peaked in 1995-96 when volatility reached very low levels. In

mid-1998, the autocorrelations dropped as volatility rose sharply. In short there is distinct cause for worry that markets were artificially smoothed during the 1995-97 periods¹.

Similarly, this study takes the view that the scam is a period of episodic volatility (event risk) which could quite easily recur. If we disregard issues of morality and legality, the scam was essentially a problem of monetary policy or credit policy. Since both the bull and bear sides of the market financed themselves through the scam in roughly equal measure, the scam was roughly neutral in terms of direct buy or sell pressure on the market. What caused a strong impact on stock prices was the vastly enhanced liquidity in the stock market. The scam was (in its impact on the stock market) essentially equivalent to monetary easing or credit expansion on a large scale. The exposure of the scam was similarly equivalent to dramatic monetary (or credit) tightening. Any sudden and sharp change in the stance of monetary policy can be expected to have an impact on the stock market very similar to the scam and its exposure. A prudent risk management system must be prepared to deal with events of this kind.

Distribution of Market Returns

The usual definition of return as the percentage change in price has a very serious problem in that it is not symmetric. For example, if the index rises from 1000 to 2000, the percentage return would be 100%, but if it falls back from 2000 to 1000, the percentage return is not -100% but only -50%. As a result, the percentage return on the negative side cannot be below -100%, while on the positive side, there is no limit on the return. The statistical implication of this is that returns are skewed in the positive direction and the use of the normal distribution becomes inappropriate.

For statistical purposes, therefore, it is convenient to define the return in logarithmic terms as $r_t = \ln(I_t/I_{t-1})$ where I_t is the index at time t . The logarithmic return can also be rewritten as $r_t = \ln(1+R_t)$ where R_t is the percentage return showing that it is essentially a logarithmic transformation of the usual return. In the reverse direction, the percentage return can be recovered from the logarithmic return by the formula, $R_t = \exp(r_t)-1$. Thus after the entire analysis is done in terms of logarithmic return, the results can be restated in terms of percentage returns.

It is worth pointing out that the percentage return and the logarithmic return are very close to each other when the return is small in magnitude. However, when there is a large return (positive or negative) the logarithmic return can be substantially different from the percentage return. For example, in the earlier illustration of the index rising from 1000 to 2000 and then dropping back to 1000, the logarithmic returns would be +69.3% and -69.3% respectively as compared to the percentage returns of +100% and -50% respectively.

¹ The cause for this high autocorrelation is a subject for further research. Some experts believe that front-running for the FIIs could have led to this phenomenon.

Distribution of Stock Market Returns




Figure 1

Figure 1 shows the probability density function of the distribution of stock market returns estimated using a gaussian kernel² with a bandwidth of 0.20 standard deviations. (In this and all subsequent density plots, the units on the X axis are in terms of the historical standard deviation calculated over the full sample). As can be seen the distribution is characterised by a thinner waist and fatter tails than the normal distribution. The summary statistics of the distribution are as follows:

Mean	0.07%
Median	0.00%
Standard Deviation (σ)	1.96%
Quartile Deviation x 0.7413 (this should equal the standard deviation for a normal distribution)	1.50%
Skewness	0.04
Excess Kurtosis (Excess of the kurtosis over the normal distribution value of 3).	5.42
Maximum	12.11% (= 6.2 σ)
Minimum	-12.54% (= -6.4 σ)

² For a description of kernel and other methods of density estimation, see Silverman (1986).

The non-normality of the distribution is evident from the large extreme values and the high excess kurtosis. However, it is well known that the principal reason for the non-normality of the unconditional distribution is that the volatility is varying over time. The observed unconditional distribution is actually a mixture of these conditional distributions of varying volatility. These conditional distributions are expected to be much closer to normality.

Modelling Time Varying Volatility

Practitioners have often dealt with time varying parameters by confining attention to the recent past and ignoring observations from the distant past. Econometricians have on the other hand developed sophisticated models of time varying volatility like the GARCH (Generalised Auto-Regressive Conditional Heteroscedasticity) model (Bollerslev, 1986).

Straddling the two are the exponentially weighted moving average (EWMA) methods popularised by J. P. Morgan's RiskMetrics® system. EWMA methods can be regarded as a variant of the practitioner's idea of using only the recent past. The practitioners' idea is essentially that of a simple moving average where the recent past gets a weight of one and data before that gets a weight of zero. The variation in EWMA is that the observations are given different weights with the most recent data getting the highest weight and the weights declining rapidly as one goes back. Effectively, therefore, EWMA is also based on the recent past. In fact, it is even more responsive than the simple moving average to sudden changes in volatility. EWMA can also be regarded as a special case of GARCH as shown below.

The simple GARCH (1,1) model can be written as follows:

$$\sigma_t^2 = \omega^2 + \beta \sigma_{t-1}^2 + \alpha r_{t-1}^2 \quad (1)$$

$r_t / \sigma_t \sim N(0, 1)$ or more generally iid with zero mean & unit variance

where r_t is the logarithmic return on day t (defined as $\ln(I_t/I_{t-1})$ where I_t is the market index on day t), σ_t is the standard deviation of r_t , α and β are parameters satisfying $0 \leq \alpha \leq 1$, $0 \leq \beta \leq 1$, $\alpha + \beta \leq 1$ and $\omega^2/(1 - \alpha - \beta)$ is the long run variance. This is the simplest GARCH model in that it contains only one lagged term each in σ and r and uses the normal distribution. More general models can be obtained by considering longer lag polynomials in σ and r and using non normal distributions.

Essentially, the GARCH model accommodates different stock market regimes by allowing the volatility of the market to vary over time. It also postulates that a large change in the index (whether positive or negative) is likely to be followed by other large changes in subsequent days. This effect is captured by using the squared return to update the estimated variance for the next day. In fact the posterior variance σ_t^2 is a weighted average of three quantities: (i) the long run variance $\omega^2/(1 - \alpha - \beta)$ with weight $(1 - \alpha - \beta)$, (ii) the prior variance σ_{t-1}^2 with weight

β and (iii) the squared return r_{t-1}^2 with weight α . The restrictions on the parameters α and β ensure that the weights are positive and sum to unity.

A special case of the GARCH model arises when $\alpha + \beta = 1$ and $\omega = 0$. In this case, it is common to use the symbol λ for β and Eq 1 takes the simpler form

$$\sigma_t^2 = \lambda \sigma_{t-1}^2 + (1 - \lambda) r_{t-1}^2 \quad (2)$$

$r_t / \sigma_t \sim N(0, 1)$ or more generally iid with zero mean & unit variance

The variance estimate can in this case be also interpreted as a weighted average of all past squared returns with the weights declining exponentially ($[1-\lambda]$, $[1-\lambda]^2$, $[1-\lambda]^3$...) as we go further and further back. Eq. 2 is therefore the Exponentially Weighted Moving Average (EWMA) model.

Initial estimation of both Eq 1 and Eq 2 using the normal distribution indicated significant non normality. They were therefore estimated using the Generalised Error Distribution (GED) which was popularised in financial econometrics by Nelson (1991). The GED with zero mean and unit variance and tail parameter ν ($0 < \nu$) is defined by the density:

$$g(x) = \frac{\nu \exp\left(-\frac{1}{2} \left| \frac{x}{\kappa} \right|^\nu\right)}{\kappa 2^{\frac{\nu+1}{\nu}} \Gamma\left(\frac{1}{\nu}\right)} \quad (3)$$

$$\kappa = \sqrt{\frac{2^{-\frac{2}{\nu}} \Gamma\left(\frac{1}{\nu}\right)}{\Gamma\left(\frac{3}{\nu}\right)}}$$

The normal distribution is the special case where $\nu = 2$. Low values of ν imply fatter tails than the normal while higher values imply thinner tails. In the GARCH estimates, the tail parameter ν of the GED was about 1.46 (as against the value of 2 for the normal distribution) implying fatter tails than the normal distribution.

Empirical Results

Model Estimation

GARCH-GED Model: Estimation of the GARCH model (Eq 1) with GED residuals produced estimates of α and β of 0.100 and 0.886 respectively implying $\alpha + \beta$ equal to 0.986 while ν was estimated to be 1.46. The estimate of ω^2 was 6.14E-06 implying a long run *standard deviation* of daily market return of about 2.1% which is close to the historical value. The log

likelihood for this model was 4735. The log likelihood ratio test rejected the hypothesis of normality ($\nu = 2$) very strongly (χ^2 with 1 df = 52.9, $P < 0.001\%$).

EWMA-GED Model: Since $\alpha + \beta$ was close to 1, the model was re-estimated after imposing the restriction that $\omega = 0$ thereby collapsing the GARCH model to the EWMA model (Eq 2) with GED residuals. The estimate of α (or λ) was 0.923 and ν was estimated to be 1.46. The log likelihood for this model was 4722. Therefore, the log likelihood ratio test rejected this model in favour of model 2 above (χ^2 with 1 df = 25.8, $P < 0.001\%$). In this model also, the log likelihood ratio test rejected the hypothesis of normality ($\nu = 2$) very strongly (χ^2 with 1 df = 61.5, $P < 0.001\%$).

EWMA-RM Model: Since the value of λ in the EWMA-GED model is fairly close to the value of 0.94 used in the RiskMetrics® methodology of J. P. Morgan, the log likelihood was computed for this model also and found to be 4720. The likelihood ratio test was unable to reject this model as against the EWMA-GED model (χ^2 with 1 df = 3.28, $P = 7\%$). In this model also, the log likelihood ratio test rejected the hypothesis of normality ($\nu = 2$) very strongly (χ^2 with 1 df = 62.7, $P < 0.001\%$).

Goodness of fit

The performance of the GARCH-GED and the EWMA-GED models can be measured by examining the distribution of the standardised residuals r_t/σ_t . The table below compares the summary statistics of the distribution of the standardised residuals with those of the normal distribution and GED. Figure 2 plots the density (tails) of the Garch-GED residuals (Gaussian kernel estimate with a bandwidth of 0.20 standard deviations) and compares it with the tails of the normal and GED densities. Figure 3 provides similar plots for the EWMA-RM residuals. These plots show mild departures from the GED for the GARCH-GED model and slightly more pronounced departures for the EWMA-RM model (observe the small hump between 3 and 4 standard deviations).

	Expected Value based on		GARCH - GED	EWMA - GED	EWMA-RM
	Normal	GED ($\nu = 1.46$)			
Standard Deviation	1.00	1.00	1.00	1.06	1.05
Quartile Deviation	1.35	1.20	1.24	1.29	1.27
Skewness	0.00	0.00	-0.01	0.01	0.03
Excess Kurtosis	0.00	0.85	1.94	1.78	1.76
Maximum			4.70	4.27	4.33
Minimum			-5.90	-6.42	-6.32
Number beyond ± 5	≈ 0	0.10	2	2	2

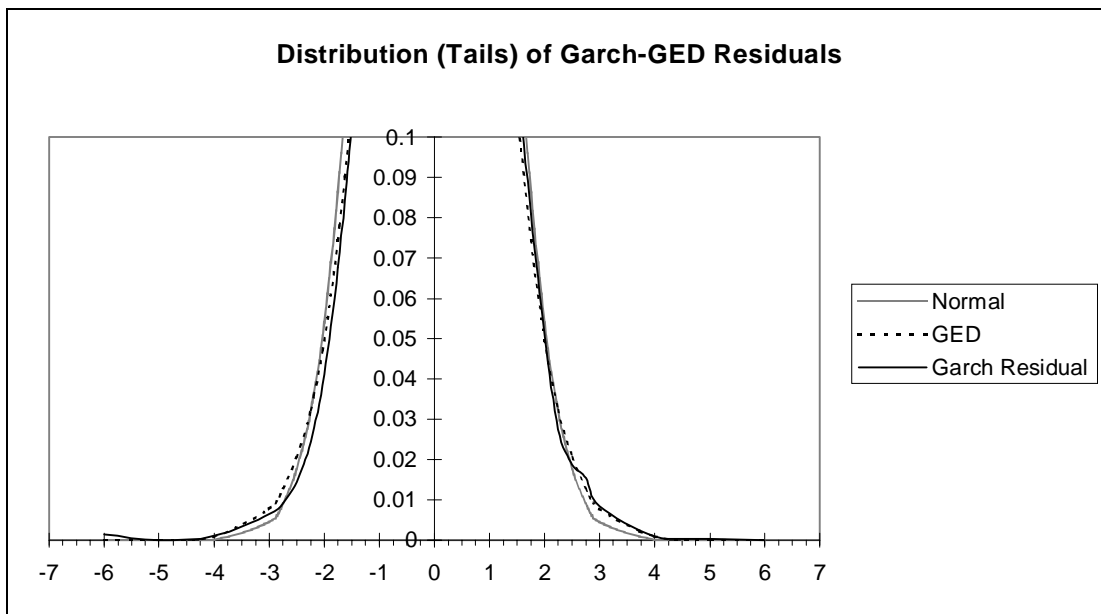


Figure 2

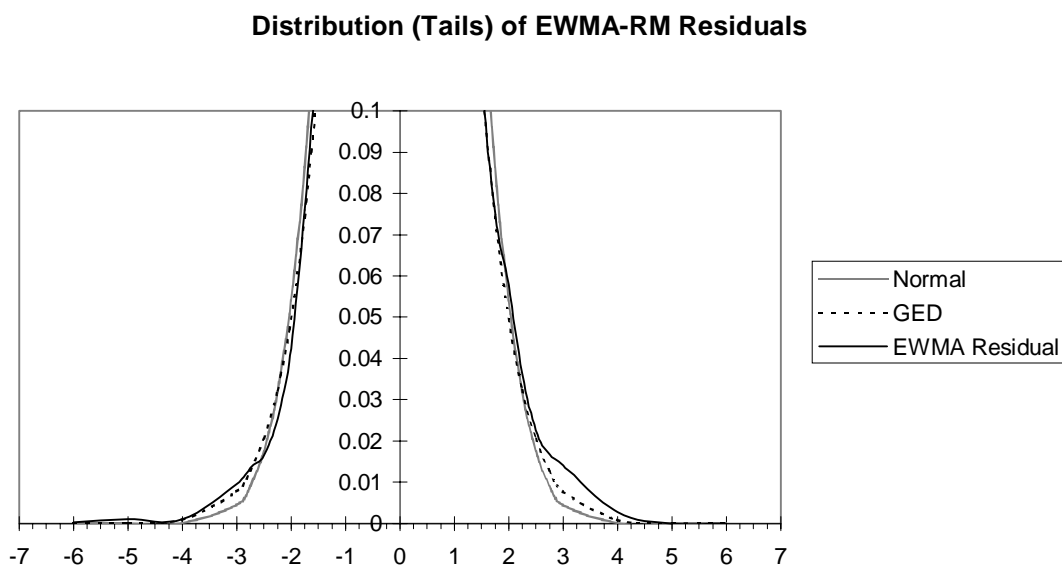


Figure 3

The standardised residuals from both the EWMA-GED and the GARCH-GED models have somewhat thinner waists and fatter tails than the GED distribution with $\nu = 1.46$. Apart from this however, all three models appear to provide reasonably good fits to the data in terms of the above broad parameters.

Value at risk

Many risk management models emphasise the calculation of Value at Risk (VaR) which is defined in terms of the percentiles of the distribution of asset values. Where the asset is an open position in index futures or in diversified stock portfolios, the VaR computation reduces to the calculation of the percentiles of the distribution of market returns. Most VaR computations are based on the 1st, 5th and 10th percentiles. The table below shows the performance of the EWMA-GED and GARCH-GED models at these risk levels (the performance of EWMA-RM is very similar to that of EWMA-GED). It is seen that the GARCH-GED model does well at all risk levels while the EWMA models do well at the 10% and 5% levels but break down at the 1% risk level.

	10% level two sided	5% level two sided	1% level two sided
Percentile as number of standard deviations ³ (using GED with $\nu = 1.46$)	1.65	2.04	2.85
Expected number of violations of VaR limit	178	89	18
EWMA-GED: Actual number of violations	183	100	31
EWMA-GED: Actual percentile	10.31%	5.63%	1.75%
EWMA-GED: Significance test of actual versus expected	Not significant	Not significant	Significant ($P \approx 0.17\%$)
GARCH-GED: Actual number of violations	161	72	20
GARCH-GED: Actual percentile	9.07%	4.06%	1.13%
GARCH-GED: Significance test of actual versus expected	Not significant	Mildly Significant ($P \approx 3.90\%$)	Not significant

At even lower risk levels like 0.50% and 0.25%, the GARCH-GED model continues to do well while the EWMA models fares poorly as shown below:

³ The cdf of the GED has to be calculated by numerical integration of the GED density. The percentile can then be obtained by a simple one dimensional search procedure or by Newton-Raphson iterations.

	0.50% level two sided	0.25% level two sided
Percentile as number of standard deviations (using GED with $\nu = 1.46$)	3.18	3.49
Expected number of violations of VaR limit	9	4
EWMA-GED: Actual number of violations	15	11
EWMA-GED: Actual percentile	0.85%	0.62%
EWMA-GED: Significance test of actual versus expected	Mildly Significant (P = 1.95%)	Significant (P = 0.20%)
GARCH-GED: Actual number of violations	13	7
GARCH-GED: Actual percentile	0.73%	0.39%
GARCH-GED: Significance test of actual versus expected	Not significant	Not significant

Salvaging EWMA Models

The above results show that though the estimated GARCH model is rather close to the EWMA models, the difference is statistically highly significant. Moreover, the EWMA models are less successful than full-fledged GARCH models in value at risk assessments at low risk levels.

Nevertheless, EWMA models have several major advantages which make them attractive for practical use:

1. EWMA involves nothing more complicated than a moving average which all market participants are familiar with. Therefore the model has the advantage of simplicity and ease of understanding.
2. EWMA models also have far greater tractability when extended to the multivariate case since the same exponential moving average technique can be applied as easily to correlations as to variances. Multivariate GARCH models by contrast are more complex and computationally more demanding.
3. GARCH models involve an estimate of the long run volatility. When the market undergoes structural changes, this long run volatility can also change. Since EWMA models do not involve any notion of a long run volatility at all, they are more robust under regime shifts.

It is therefore worthwhile to see whether it is possible to salvage the EWMA models by making some suitable adjustments. The starting point for any such adjustment is the distribution of the standardised residuals r_t/σ_t which we have already examined earlier. The plots of the density of the standardised residuals showed only mild departures from the GED for the GARCH-GED model (Figure 2). In case of the EWMA model (Figure 3) we observed

a small hump in the density between 3 and 4 standard deviations. It turns out that for value at risk purposes at most common risk levels, the mild departures of the GARCH-GED residuals from the GED are unimportant. However, in case of the EWMA, the hump in the density becomes important at low risk levels.

The clue to salvaging the EWMA models lies here. The hump in the density plot must be taken into account while using EWMA models for value at risk purposes. For example, while the GED suggests that 1% value at risk estimates can be obtained by using 2.85 standard deviations, the hump suggests that we must use a slightly higher value - say 3 standard deviations. In fact, use of 3 standard deviations is a normal rule of thumb for distributions with a moderate degree of non normality.

Using the 3 standard deviations rule, we find that the 1% VAR limit was crossed 22 times as against the expected number of 18 violations. The hypothesis that the true probability of a violation is 1% cannot be rejected at even the 5% level of statistical significance though we have a sample size of over 1750. The actual number of violations is therefore well within the allowable limits of sampling error. In the terminology of the Bank for International Settlements⁴, these numbers are well within the “Green Zone” where the “test results are consistent with an accurate model, and the probability of accepting an inaccurate model is low”.

Margins

Based on a similar analysis, the author recommended⁵ that margins levied by the derivatives exchanges for the proposed index futures contracts should be based on the EWMA-RM model. Since the volatility estimates are for the logarithmic return, the $\pm 3 \sigma$ limits for a 99% VAR would specify the maximum/minimum limits on the logarithmic returns not the percentage returns. To convert these into percentage margins, the logarithmic returns would have to be converted into percentage price changes by reversing the logarithmic transformation. Therefore the percentage margin on short positions would be equal to $100(\exp(3\sigma)-1)$ and the percentage margin on long positions would be equal to $100(1-\exp(-3\sigma))$. This implies slightly larger margins on short positions than on long positions, but the difference is not significant except during periods of high volatility where the difference merely reflects the fact that the downside is limited (prices can at most fall to zero) while the upside is unlimited.

⁴ *Supervisory framework for the use of ‘backtesting’ in conjunction with the internal models approach to marker risk capital requirements*, Basle Committee on Banking Supervision, January 1996

⁵ Varma, J. R., Chairman (1999), *Report of the Committee on Risk Containment in the Derivatives Markets*, Securities and Exchange Board of India, Mumbai.

The market movements, margins and margin shortfalls based on this approach are shown graphically in Figure 4. The summary statistics about the actual margins on the sell side are tabulated below.

Sell Side Margins							
<i>Summary</i>			<i>Frequency Distribution</i>				
<i>Average</i>	<i>Max</i>	<i>Min</i>	< 5%	5 - 10%	10 - 15%	15 - 20%	> 20%
5.49%	21.73%	2.04%	52.51%	41.75%	4.17%	1.35%	0.23%

A Closer Look at VAR Violations

Taking a closer look at the actual violations (See Figure 5), it is seen that most of the violations take place when the market move is large and the violation is typically a small fraction of the market move. This implies that in most cases, the model is able to correctly forecast that the markets are in a volatile period and step up the margins accordingly to protect market integrity.

There are only two exceptions to this pattern. The first exception is March 31, 1997 when the sudden withdrawal of support to the then government by the major supporting party led to a sharp fall in the market. This is the kind of event risk which a statistical model cannot predict and against which the only protection can be a second line of defence (broker net worth). The second exception is October 28, 1997 when the global equity meltdown triggered by sharp falls in the Asian markets and in the US market drove the Indian market also down.

It is conceivable, though by no means certain, that more sophisticated statistical models which can estimate volatility contagion across several financial markets could have provided better protection against the market drop of October 28, 1997. The development of multivariate models of volatility estimation that can account for such contagion is a topic for further research. Practical utility of such a model would however be contingent on the ability of the derivatives exchange / clearing corporation to make a margin call shortly before the market opens in Mumbai based on the market movement in New York (previous day close), Tokyo (same day close) and Hong Kong (same day morning session).

Nifty returns plotted against confidence limits of 3 standard deviations




Figure 4

Hypothetical Margin Shortfalls in Nifty Using 3 Standard Deviations




Figure 5

References

Bank for International Settlements (1996), *Supervisory framework for the use of 'backtesting' in conjunction with the internal models approach to market risk capital requirements*, Basle Committee on Banking Supervision, January 1996

Varma, J. R., Chairman (1999), *Report of the Committee on Risk Containment in the Derivatives Markets*, Securities and Exchange Board of India, Mumbai.