Characterisation of the tail behaviour of financial returns: an empirical study from India's stock market

(Preliminary draft)

Mandira Sarma*

25 February, 2004

Abstract

This paper uses extreme value theory to explicitly model the tail regions of the innovation distribution of the return series of S & P CNX Nifty, the prominent Indian equity index. We model each tail separately by fitting a Generalised Pareto distribution to the observations lying beyond certain threshold that marks the beginning of the tail region. In line with the much discussed stylised feature of financial returns, we find existence of tail-thickness in both the lower and the upper tails of the marginal distribution of Nifty logarithmic returns. However, we do not find any evidence of asymmetric tails.

KEY WORDS: Extreme value theory, tail behaviour, Peaks-over-threshold model

JEL Classification: C10, C13, C22, G10

*EURANDOM, P.O.Box 513, 5600 MB Eindhoven, The Netherlands. Email:sarma@eurandom.tue.nl

Characterisation of the tail behaviour of financial returns: an empirical study from India's stock market

1 Introduction

The analysis of tail behaviour of asset returns is important from the point of view of risk management, as financial risk management is all about understanding large movements of asset prices. Explicit forms of the tails of the distribution provide important information to risk managers and investors.

Empirical studies have established that the distribution of speculative asset returns tend to have heavier tails than the Gaussian distribution tails (Mandelbrot, 1963; Pagan, 1996). Furthermore, very often such distributions are found to have asymmetric tails. Such stylised features of financial returns provides interesting insight into the economics of financial markets and calls for appropriate methodologies of modelling such behaviour.

Conditional heteroskedasticity models of Engle (1982) and Bollerslev (1986) and their various modifications do incorporate some of these stylised features which emanate due to phenomenon such as volatility clustering in financial data. Although the conditional heteroskedasticity models can explain part of the non-Gaussina features of the unconditional distribution, it is often found that features like heavy tails may persist even after accounting for conditional heteroskedasticity.

This paper uses the recent developments of extreme value theory to empirically characterise the tails of the unconditional distribution of the prominent Indian stock market index, the S&P CNX Nifty. Using the "Peaks-Over-Threshold" (POT) model (McNeil and Frey, 2000) we estimate each tail separately by fitting a Generalised Pareto distribution to the observations lying beyond certain threshold that marks the beginning of the tail region.

In line with the much discussed stylised feature of financial returns, we find existence of tail-thickness in both the lower and the upper tails of the unconditional distribution of Nifty logarithmic returns. However, we do not find asymmetric tail behaviour.

The rest of this paper is organised as follows. Section 2 describes the methodology used in this paper, along with an overview of extreme value theory including the main approach of tail estimation, the "Peaks-Over-Threshold (POT)" model. Section 3 provides an account of the empirical analysis of the tails of the Nifty innovations and presents the important results. Section 4 concludes this paper.

2 Methodology

Suppose that the asset return r_t at time t can be described by the following time series specification

$$r_t = \mu_t + \sigma_t z_t \tag{1}$$

where μ_t describes the time varying mean, σ_t is the time varying volatility dynamics and z_t's are iid white noise innovations.

The distribution of z_t, particularly its tail regions, is the focus of this paper. We use the "Peaks-Over-Threshold" (POT) model (McNeil and Frey, 2000) to estimate the tails regions of the innovation distribution. The POT model is based on the "Pickands-Balkema-de Haan theorem" which postulates that the distribution of the observations in excesses of certain high threshold can be approximated by a Generalised Pareto distribution. Application of this theorem requires z_t to be *iid* and therefore it is crucial to have appropriate specifications of μ_t and σ_t such that z_t is purely white noise and does not contain any time dependence.

In the following subsection we briefly describe the relevant concepts from extreme value theory and then discuss in details the Peaks-Over-Threshold model.

2.1 An overview of Extreme Value Theory

The classical Extreme Value theory (EVT) deals with the study of the asymptotic behaviour of extreme observations (maxima or minima of n random realisations) of random variables.

Suppose that $X \in (l, u)$ is a random variable with density f and cdf F. Let $X_1, X_2, ...X_n$ be n independent realisations of the random variable X. Define the extreme observations as

$$Y_n = max\{X_1, X_2, ..., X_n\}$$

$$Z_n = min\{X_1, X_2, ..., X_n\}$$

Extreme value theory deals with the distributional properties of Y_n and Z_n as n becomes large.

It can be easily shown that the exact distributions of the extreme observation is degenerate in the limit. In order to find a distribution of interest which is non-degenerate, the extrema Y_n and Z_n are transformed with a scale parameter $a_n (> 0)$ and a location parameter $b_n \in R$, such that the distribution of the standardised extrema

$$\frac{Y_n - a_n}{b_n}$$

$$\text{and} \quad \frac{Z_n - a_n}{b_n}$$

is non-degenerate.

The two extremes, the maximum and the minimum are related by the following relation:

$$min\{X_1, X_2, ..., X_n\} = -max\{-X_1, -X_2, ..., -X_n\}$$

Therefore, all the results for the distribution of maxima leads to an analogous result for the distribution of minima and vice versa [1].

[1] If $max\{X_1, X_2, ...X_n\} \sim H(x)$ then $min\{X_1, X_2, ...X_n\} \sim 1 - H(-x)$.

2.2 The Fisher-Tippett Theorem

The Fisher-Tippett theorem (1928) is a fundamental result in EVT. The importance of this result is that it exhibits the possible limiting forms for the distribution of Y_n under linear transformations even without the exact knowledge of the underlying distribution F. The "Fisher-Tippett Theorem", also known as the "Extremal type theorem" states thus:

If $\exists$ constants $a_n(> 0)$ and $b_n \in R$ such that

$$\frac{Y_n - a_n}{b_n} \xrightarrow{d} H \quad \text{as} \;\; n \to \infty$$

for some non-degenerate distribution H, then H must be one of the only three possible 'extreme value distributions', namely, the Gumbel, the Fréchet and the Weibull distributions.

In that case, X (and the underlying distribution F) is said to belong to the (maximum) domain of attraction of the extreme value distribution H. It is denoted by $X \in DA(H)$.

2.3 The Generalized Extreme Value Distribution

The three families of extreme value distributions, viz. the Gumbel, the Fréchet and the Weibull, can be nested into a single parametric representation, as shown by Jenkinson and Von Mises. This representation is known as the "Generalised Extreme Value" (GEV) distribution, and given by

$$H_\xi(x) = exp\{-(1 + \xi x)^{-\frac{1}{\xi}}\} \tag{2}$$

where

$$1 + \xi x > 0$$

The support of ξ is

$$x \;\; > \;\; -\frac{1}{\xi} \;\; \text{if} \;\; \xi > 0$$

$$
\begin{aligned}
x &< \frac{1}{\xi} \text{ if } \xi < 0 \\
x &\in R \text{ if } \xi = 0
\end{aligned}
$$

The parameter ξ, called the tail index, determines the tail-thickness. When $\xi > 0$, we get the Fréchet distribution family, which incorporates in it the fat-tailed distributions such as Student's-t or the Stable Paretian distributions. The marginal distribution of a stationary GARCH process is also in the domain of attraction of the Fréchet family. The case when $\xi = 0$ is the case of the Gumbel class, which describe the thin-tailed distributions like the normal or log-normal distribution. Finally, when $\xi < 0$ we get the Weibull distribution which describe distributions without a tail, but a finite end-point, such as the uniform and the beta distribution.

2.4 The Pickands-Balkema-de Haan Theorem

Suppose that $X_1, X_2, ...X_n$ are n independent realisations of a random variable X with a distribution function $F(x)$. Let u be the finite or infinite right endpoint of the distribution F. Let the distribution function of the excesses over certain (high) threshold k is denoted by $\Phi_k(x)$.

The Pickands-Balkema-de Haan theorem (Balkema & de Haan 1974; Pickands 1975) states that if the distribution function $F \in DA(H_\xi)$ then $\exists$ a positive measurable function $\beta(k)$ such that

$$
\lim_{k \to u} \sup_{0 \le x < u-k} |\Phi_k(x) - G_{\xi,\beta(k)}(x)| = 0
$$

and vice versa, where $G_{\xi,\beta(k)}(x)$ denote the Generalised Pareto distribution.

The above theorem states that as the threshold k becomes large, the distribution of the excesses over the threshold tends to the Generalised Pareto distribution, provided the underlying distribution F satisfies the extremal-types theorem [2].

[2] Almost all common continuous distributions used in finance do satisfy the extremal types theorem; hence the Pickands-Balkema-de Haan theorem is fairly general.

2.5 The Generalised Pareto Distribution (GPD)

The GPD is given by

$$G_{\xi,\beta}(x) = \begin{cases} 1 - (1 + \xi x/\beta)^{-1/\xi}\ ; & if\ \xi \neq 0 \\ 1 - exp(-x/\beta); & if\ \xi = 0 \end{cases} \tag{3}$$

where $\beta > 0$, and the support of x is $x \geq 0$ when $\xi \geq 0$ and $0 \leq x \leq -\beta/\xi$ when $\xi < 0$.

In the POT model, a certain threshold is identified to define the starting of the tail of the return distribution. Then the distribution of the 'excesses' over the threshold point is estimated. There are two approaches of estimating the 'excess' distribution, viz. the semi-parametric models based on the Hill estimator (Danielsson and de Vries, 2000) and the fully parametric model based on the Generalised Pareto distribution (GPD) (McNeil and Frey, 2000). The Hill estimator based approach is limited in its application as it requires the assumption of fat tails of the underlying return distribution. On the other hand, the GPD version is applicable to any kind of distribution, fat-tailed or not. This approach utilises the Pickands-Balkema-de Haan theorem to fit a generalised Pareto distribution to the excesses over specific thresholds.

The following sections describe the GPD approach of the POT model in details.

2.6 The Peaks-over-Threshold Model: the GPD approach

The POT model provides for a framework of estimating the tails (positive or negative tails) of the marginal distribution of the returns by estimating what is known as the distribution of *excesses* over certain threshold point which identifies the starting of the tail.

The distribution of *excesses* over a high threshold k on the portfolio's loss distribution F is defined by

$$\Phi_k(y) = Pr\{X - k \leq y | X > k\}$$

In terms of the underlying loss distribution F,

$$\Phi_k(y) = \frac{F(y+k) - F(k)}{1 - F(k)} \tag{4}$$

Pickands-Balkema-de Haan theorem states

$$\Phi_k(y) \to G_{\xi,\beta(k)}(y) \tag{5}$$

for

$$k \to u$$

Thus, using the Pickands-Balkema-de Haan theorem, one can model the distribution of the excesses over the threshold k as a GPD, provided the threshold is sufficiently high.

Setting $x = k + y$ and using (4) and (5), we can rewrite F as

$$F(x) = (1 - F(k))G_{\xi,\beta}(x - k) + F(k) \tag{6}$$

for $x > k$.

Using HS estimate for $F(k)$ and ML estimates of the GPD parameters gives rise to the following tail estimator formula

$$\hat{F}(x) = 1 - \frac{N_k}{N}\left(1 + \hat{\xi}\frac{x-k}{\hat{\beta}}\right)^{-\frac{1}{\hat{\xi}}} \tag{7}$$

For a given probability $p > F(k)$, a tail quantile is estimated by inverting the tail estimator formula (7),

$$\hat{q}_p = k + \frac{\hat{\beta}}{\hat{\xi}}\left(\frac{N}{N_k}(1-p)^{-\hat{\xi}} - 1\right) \tag{8}$$

Equations (7) and (8) provide the basic formulae to estimate the tail probabilities and the tail quantiles.

3 An empirical characterisation of the Nifty tails

In this section, we present a POT analysis of the innovation distribution of the returns of the S & P CNX Nifty (or in short, Nifty) index of India's National Stock Exchange.

3.1 Data

The data consists of 3172 daily percentage logarithmic returns of the Nifty index from 3 July 1990 till 13 January 2004[3]. Some descriptive statistics of the data can be found in the first column of Table 2.

3.2 Estimation and results

3.2.1 Time series analysis

Figure 1 depict the correlograms of the returns and the squared returns of the Nifty logarithmic returns. As seen from this figure, there is significant autocorrelation in the returns as well as the suqred returns.

It is crucial to have an accurate specification of the time series dynamics of the returns in order to obtain an *iid* residual series to which one can apply the POT model.

A time series specification search in terms of SBC criterion for the mean dynamics of the logarithmic Nifty returns suggests a AR(1) model to be the most suitable description for the mean equation. As far as the volatility dynamics is concerned, we have chosen a GARCH(1,1) model, based on the SBC criterion[4].

Table 1 presents the estimated parameters of the mean and volatility equations of the Nifty returns. All the parameters – the constants and the AR parameter in the mean equation, and

[3] The data has been obtained from the web page of National Stock Exchange of India, ¡www.nseindia.com¿.

[4] Values of the AIC and SBC criteria for various specifications of the time series can be obtained on request.

the constant, the ARCH parameter and the GARCH parameter in the volatility equation, are found to be significant.

3.2.2 The standard residuals

Figure 2 depict the correlogram of the standard residuals from the AR(1)-GARCH(1,1) specification of the logarithmic returns of Nifty. It is evident from the correlograms that while there is significant autocorrelation in the returns and the squared returns, the residual series is devoid of time dependence, and therefore may be considered *iid*.

Table 2 presents the values of the first four unconditional moments of the Nifty series and the standard residuals obtained from an AR(1)-GARCH(1,1) specification of the Nifty logarithmic returns. The values of these descriptive statistics indicate the existence of leptokurtosis in both the raw returns as well as the standard residuals. Thus, incorporation of conditional heteroskedasticity in the form of a GARCH(1,1) specification has only partially removed the leptokurtosis in the unconditional distribution of Nifty returns. A substantial part of the leptokurtosis still exists in the unconditional distribution even after the time varying volatility has been removed.

Table 3 presents in Panel A and Panel B, the values of the test statistics for testing the significance of skewness, excess kurtosis and autocorrelation (upto lag 35) along with their respective p-values for the original return series and the residual series. These results indicate that the return series has significant excess kurtosis and autocorrelation. The residual series is found to have significant excess kurtosis but it does not possess significant autocorrelation. In both the returns and the standard residuals, the skewness coefficient is found to be insignificant, implying distributional symmetry. Existence of highly significant excess kurtosis in both the returns and the residual series imply that the distribution of the returns as well as the residuals is far from being normal, instead, they are heavy tailed. As table 3 indicates, the residual series is found to be free from autocorrelation and hence can be considered *iid*.

A non-parametric Anderson-Darling test was performed on the standard residuals to test whether or not the residual series may be considered to be normally distributed. The results of this test, as shown in Panel C of Table 3, indicate an outright rejection of normality of the residual series.

3.3 Modeling peaks-over-thresholds

The Pickands-Balkema-de Haan theorem offers the generalised Pareto distribution as a natural choice for the distribution of excesses (peaks) over sufficiently high thresholds. However, while choosing an appropriate threshold, one faces an unpleasant trade off between bias and variance. The theoretical consideration suggests that the threshold should be as high as possible for the Pickands-Balkema-de Haan theorem to hold good, but in practice, too high a threshold might leave us with very few observations above the threshold for estimating the GPD parameters, leading to statistical imprecision and very high variance of the estimates[5].

There is no correct choice of the threshold level. While McNeil and Frey (2000), McNeil (1997) and McNeil (1999) use the "mean-excess-plot" as a tool for choosing the optimal threshold level[6], Gavin (2000) uses an arbitrary threshold level of 90% confidence level (i.e. the largest 10% of the positive and negative returns are considered as the extreme observations).

In this paper we follow Neftci (2000) and choose the threshold level as 1.645 times the unconditional variance of the residuals[7]. On the both sides of the tails, observations lying beyond 1.645 times the unconditional variance are considered to be extremes. After identifying the thresholds thus, we use the observations in excess of the threshold to obtain the maximum likelihood estimates of the GPD parameters $\hat{\xi}$ and $\hat{\beta}$.

Table 4 presents the estimated threshold point, the number of extreme observations beyond the threshold, and the results of the maximum likelihood estimation of the GPD to the excesses (peaks) over the chosen threshold for the lower tail and the upper tail respectively of the i.i.d. residual series obtained by fitting an AR(1)-GARCH(1,1) model to the Nifty

[5] For a discussion on this issue, see McNeil and Frey (2000).

[6] Details on mean-excess-plots can be found in McNeil and Frey (2000) and Embrechts et al. (1997).

[7] This represents the 5% of extreme observations if the data were normally distributed. We tried with the mean excess plots but did not get a well behaved linear mean excess plot.

returns. The first column of this table gives the threshold points corresponding to the lower tail (-1.6460) and the upper tail (1.6460).

The second column of table 4 gives the number of extreme observations beyond the thresholds on both the tails, the third column gives the estimated cdf at the thresholds, the fourth and the fifth columns presents the ML estimation of the GPD parameters (with SEs in parenthesis) fitted to the observations in excess over the thresholds.

As seen from table 4, the estimated $\hat{\xi}$ is found to be positive for both the lower tail and the upper tail of the innovation distribution. This indicate that the tails on both sides of the innovation distribution are heavy.

The estimates of the GPD parameters can be used in the tail quantile estimation formula (8) to estimate the tails of the distribution.

3.4 Estimating the tails

Table 5 presents some of the estimated quantiles on the lower tail and the upper tail along with the empirical quantiles and the corresponding quantiles on the standard normal curve. These quantiles are estimated for a fixed value of probability level, and hence they may be considered as the unconditional Value-at-Risk measures for the i.i.d. residual series obtained from three models – the EVT-based GPD approach, the historical simulation and the Normal distribution model. The first column of this table indicates the probability levels and the second, third and the fourth columns give the corresponding quantiles. The quantiles on the second column are estimated using the GPD fit presented in Table 4. The quantiles on the third column are empirically observed quantiles and the ones on the fourth column are the corresponding quantiles on the standard normal curve. Panel A of Table 5 provide the quantiles on the lower tail and the Panel B presents the quantiles on the upper tail.

Table 5 shows that the estimated tail quantiles are closer to the empirical quantiles than that of a normal distribution approximation. This implies that a normal distribution approximation of the underlying DGP would lead to misleading risk estimates.

Figure 3 present a graph of the estimated cdf for the lower and the upper tails compared with empirically observed tails and the corresponding normal approximation. As can be seen from this graphs, the estimated cdf fits the empirical cdf much better than the normal approximation. While normal approximation fails to capture the tail behaviour of the data on the upper tail, the POT model is able to capture both the tails precisely.

3.5 The KS test for discrepancy

To test for the significance of difference between the estimated and empirical tails, we carry out a non-parametric Kolmogorov-Smirnov (KS) test, as described below.

Suppose that $F(x)$ and $G(x)$ denote respectively the empirical and the estimated distribution functions.

We test the following hypothesis:

$$H_0 : F(x) = G(x) \tag{9}$$

against the alternative hypothesis

$$H_1 : F(x) \neq G(x) \tag{10}$$

This hypothesis tests whether the estimated cdf is significantly different from the empirical cdf. The Kolmogorov-Smirnov (KS) statistic to test the hypothesis is as follows.

$$D = \sup_x |F(x) - G(x)|$$

Table 6 provides the estimated KS-statistics for testing the null hypothesis (9) against the alternative hypothesis (10), for the lower and the upper tails. The discrepancy between the estimated and empirical tails is found to be insignificant at 0.05 level of significance, thus indicating that the estimated tails are not significantly different from the empirically observed tails at 0.05 level.

4 Summary and concluding remark

In this paper we carry out an empirical analysis of the tail regions of the innovation distribution of the returns of India's Nifty index. By using the recent developments in Extreme value theory, we have estimated the explicit forms of the lower and the upper tails of the innovation distribution of Nifty log returns. In line with the much discussed stylised phenomenon of heavy tailed behaviour, our analysis finds that the innovation distribution of Nifty logarithmic returns are heavy tailed. However, we do not find any evidence of asymmatric tail behaviour in the Nifty innovations. The explicit forms of the tails provide important insight to investors and risk managers. In particular, the results show that the use of conventional methodologies such as the normal distribution model to estimate the tail-related risk measures may lead to faulty estimation of risk.

References

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

Danielsson J, de Vries CG, 2000. Value-at-Risk and extreme returns. In *Extremes and Integrated Risk Management*, edited by Embrechts P, 85–106. Risk Books.

Embrechts P, Kluppelberg C, Mikosch T, 1997. *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag Berlin Heidelberg.

Engle R, 1982 July. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* **50**:987–1007.

Gavin J, 2000. Extreme value theory- an empirical analysis of equity risk. UBS Warburg working paper.

Mandelbrot BB, 1963. The variation of certain speculative prices. *Journal of Business* **36**:392 – 417.

McNeil AJ, 1997. Estimating the tails of loss severity distributions using extreme value theory. *ASTIN Bulletin* **27**:117–137.

McNeil AJ, 1999. Extreme value theory for risk managers. In *Internal Modelling and CAD II*, 93–113. Risk Books.

McNeil AJ, Frey R, 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* **7**:271–300.

Neftci SN, 2000. Value at Risk Calculations, Extreme Events, and Tail Estimation. *The Journal of Derivatives* **Spring**:1–15.

Pagan A, 1996. The econometrics of financial markets. *Journal of Empirical Finance* **3**:15 – 102.

Table 1 Estimation of AR(1)-GARCH(1,1) model

This table presents the estimated parameters, their SE's and the corresponding confidence intervals of the AR(1)-GARCH(1,1) model fitted to the Nifty logarithmic returns.

Parameter	Estimates	SE	Confidence bounds
The mean equation:			
Constant	0.057	0.025	(0.007, 0.107)
AR	0.149	0.019	(0.112,0.187)
The variance equation:			
Constant	0.072	0.015	(0.044, 0.101)
ARCH	0.106	0.012	(0.083, 0.129)
GARCH	0.874	0.013	(0.849, 0.899)

Table 2 Descriptive statistics: Returns and Standard residuals

This table presents some descriptive statistics of the logarithmic Nifty return series and the standard residuals extracted from a AR(1)-GARCH(1,1) specification of the Nifty returns.

	Minimum	Maximum	Mean	Variance	Skewness	Kurtosis
Returns	-12.5219	12.0861	0.06151	3.2542	0.0166	8.1138
Standard residuals	-6.0434	6.4888	-0.0063	1.0014	0.0555	5.4444

Table 3 Tests for skewness, kurtosis, auto correlation and normality

This table presents the values of the test statistics and the corresponding p-values of the tests of skewness, kurtosis and autocorrelation for the raw data (Panel A) and the standard residuals (Panel B). Panel C presents the results of the non-parametric Anderson-Darling test of normality of the residual series.

	statistic	p-value
Panel A: The returns series (r_t)		
skewness	2.3385	0.126
kurtosis	953.4722	0.000
H.C. Ljung-Box	55.3965	0.016
Panel B: The residual series (z_t)		
skewness	1.2779	0.258
kurtosis	28.2650	0.000
H.C. Ljung-Box	47.1672	0.082
Panel C: Anderson-Darling (AD) test of normality		
Calcuted AD	$AD_{0.01}$	$AD_{0.05}$
5.9553169*	1.035	0.752

Table 4 Results of the GPD estimation

This table provides the results of the estimated GPD parameters fitted to the excesses over the chosen threshold. The first column gives the threshold points on both the left and the right tails corresponding to the 1.645σ level of threshold. The second column presents the number of observations beyond the threshold level and the third column gives the estimated cdf of the tails at the respective threshold points. The fourth and the fifth columns present the Pseudo-maximum-likelihood estimation of the GPD parameters fitted to the excesses over the thresholds, along with the standard errors of estimation within parenthesis.

	u	N_u	F_u	$\hat{\xi}$	$\hat{\beta}$
left tail	-1.6460	134	0.042	0.2565	0.4348
				(0.1047)	(0.0582)
Right tail	1.6460	141	0.956	0.1357	0.5563
				(0.1052)	(0.0747)

Figures in parenthesis indicate standard error

Table 5 Estimated quantiles on the i.i.d. residuals

This table provides some of the estimated quantiles on the tails along with the empirical quantiles as well as the corresponding quantiles on the standard normal distribution. Panel A deals with the lower tail and Panel B deals with the upper tail. The column 1 gives the probability level. Columns 2, 3 and 4 give the estimated, empirical and standard normal distribution quantiles corresponding to these probability levels.

p	EVT	Empirical	Normal
Panel A: Quantiles on the lower tail			
0.050	-1.5757225	-1.5548851	-1.6448536
0.040	-1.6713663	-1.6816055	-1.7506861
0.030	-1.8030331	-1.8210692	-1.8807936
0.020	-2.0059049	-1.9730586	-2.0537489
0.010	-2.4055370	-2.3120576	-2.3263479
0.009	-2.4727464	-2.3657120	-2.3656181
0.008	-2.5500619	-2.4359087	-2.4089155
0.007	-2.6405871	-2.5272382	-2.4572634
0.006	-2.7490203	-2.8313273	-2.5121443
0.005	-2.8829285	-2.9620767	-2.5758293
0.004	-3.0555745	-3.1300396	-2.6520698
0.003	-3.2932457	-3.2823680	-2.7477814
0.002	-3.6594484	-4.0362250	-2.8781617
0.001	-4.3808223	-4.4296692	-3.0902323
Panel B: Quantiles on the upper tail			
0.950	1.5815594	1.5714796	1.6448536
0.960	1.7056075	1.6896320	1.7506861
0.970	1.8711778	1.8395167	1.8807936
0.980	2.1157883	2.0993880	2.0537489
0.990	2.5664234	2.5935117	2.3263479
0.991	2.6387088	2.7360360	2.3656181
0.992	2.7207499	2.8044265	2.4089155
0.993	2.8153604	2.8591941	2.4572634
0.994	2.9267334	2.9274304	2.5121443
0.995	3.0615027	2.9872829	2.5758293
0.996	3.2310498	3.4531877	2.6520698
0.997	3.4573486	3.5061388	2.7477814
0.998	3.7916784	3.5938781	2.8781617
0.999	4.4075994	4.1014217	3.0902323

Table 6 Results of the Kolmogorov-Smirnov tests

	Upper tail	Lower tail
D	0.2794	0.1574

Critical value of D at 0.05 level of significance = 0.467

Figure 1 Correlogram of returns, squared returns and the standard residuals

This figure provides the correlograms of the Nifty returns and squared returns. The vertical lines depict the values of partial autocorrelation at various lags while the horizontal lines give the statistical bounds.

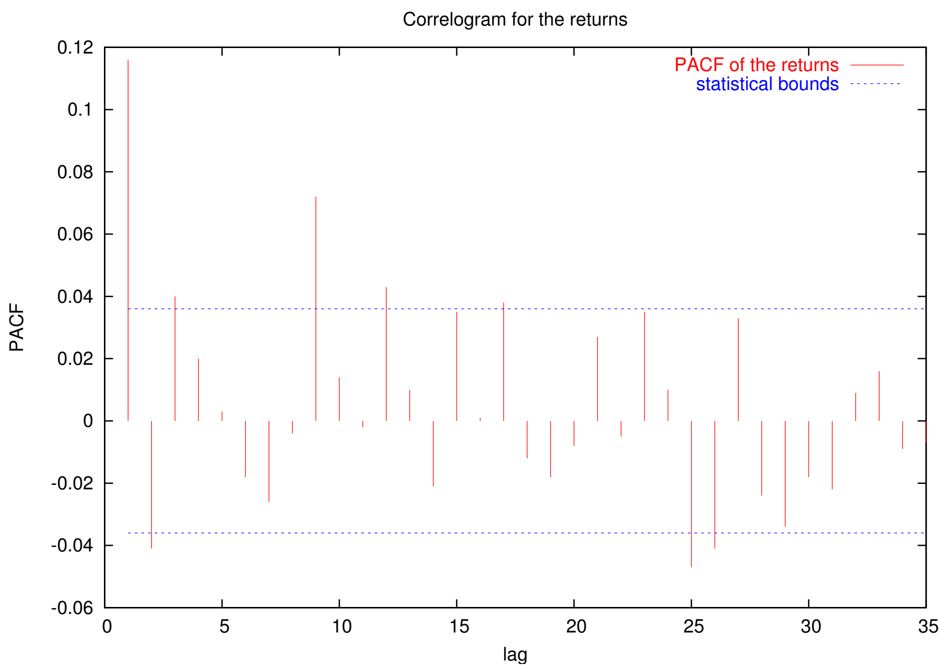

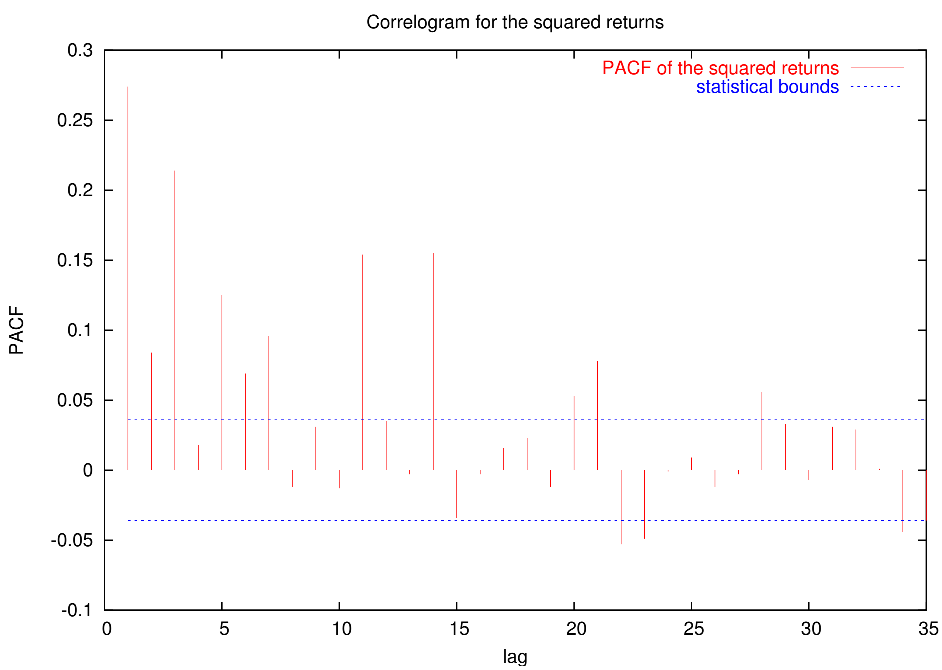

Figure 2 Correlogram of the standard residuals

This figure provides the correlogram of the standard residuals extracted from an AR(1)-GARCH(1,1) specification to the Nifty logarithmic returns. The vertical lines depict the values of partial autocorrelation at various lags while the horizontal lines give the statistical bounds.

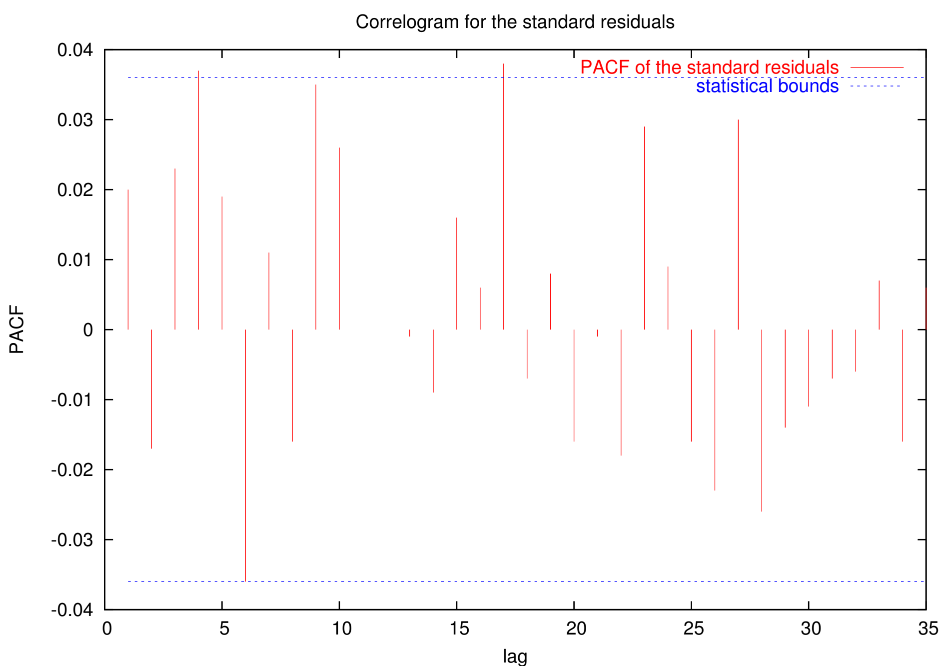

Figure 3 The tails

This figure depicts the lower tail (the graph on the top) and the upper tail (the graph on the bottom) of the innovation distribution of Nifty returns. In both the graphs, we draw estimated cdf, empirical cdf and tha normal approximations corresponding to the observed data.