Characterisation of the tail behaviour of financial returns : studies from India

Document status and date:
Published: 01/01/2005

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

Characterisation of the tail behaviour of financial returns: studies from India[*]

Mandira Sarma[†]

February 1, 2005

Abstract

In this paper we explicitly model the tail regions of the innovation distribution of two important series from the emerging financial markets of India, viz., Nifty, the equity index, and 30-day interest rate from the inter bank currency money market. Using the recent developments of extreme value theory, we estimate the tails by fitting a Generalised Pareto distribution to the observations lying beyond certain thresholds that mark the beginning of the tail regions. In line with the much discussed stylised features of financial returns, we find existence of tail-thickness, as indicated by the positive value of the tail indexes of the Generalised Pareto fit. Further, the left tail of each series is found to be heavier than the right tail. Thus, each series display asymmetric and heavy tailed behaviour.

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 Phone +31 40 247 8105 Fax +31 40 247 8190

1 Introduction

The analysis of tail behaviour of asset returns is important from the point of view of risk management, as market risk management is all about understanding large movements of asset prices. Explicit forms of the tails of asset returns distribution provide important information on the likelihood of rare but probable extreme swings in the asset prices.

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). Further, very often such distributions are found to have asymmetric tails. Such stylised features of financial returns provide 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 that emanate due to phenomenon such as volatility clustering in financial data. Although the conditional heteroskedasticity models can explain part of the non-Gaussian features of the innovation (or noise) distribution, it is often found that features like heavy tails may persist even after accounting for conditional heteroskedasticity.

This paper uses recent developments in extreme value theory to empirically characterise the tails of the innovation distributions of two important asset returns series from the emerging financial markets of India. The two series are the S & P CNX Nifty (hereafter Nifty) that represents the equities market and the 30-day money market interest rates (hereafter IR) representing the 30-day inter bank call money market. Both these series are compiled at the National Stock Exchange (NSE) of India which was set up in 1994 to encourage financial market reforms through modernisation and competition. Initiated with a state of the art market microstructure, NSE provides electronic trading platforms for fixed income securities and equity markets in India. It has brought overwhelming reforms in the financial markets in India, leading to improved transparency, reduced transaction costs and a considerable increase in the volumes traded in India's financial markets.

Nifty is a well diversified index of 50 most liquid stocks traded at NSE that cover 24 sectors of Indian economy. It was launched in April 1996 as a market capitalisation weighted index and it serves as underlying for various index-based derivatives traded at NSE. Further there are many mutual funds linked with Nifty.

The MIBID (Mumbai Inter-bank Bid Rate) and MIBOR (Mumbai Inter-bank Offer Rate) series are compiled at NSE by using a methodology of polling and bootstrapping[1]. These rates are widely accepted as benchmark interest rates corresponding to Mumbai's inter-bank call money market.

Using the "Peaks-Over-Threshold" (POT) model (McNeil and Frey, 2000) of extreme value theory, we estimate each tail (the left and the right tail) of each of the above series by fitting a Generalised Pareto distribution to the observations lying beyond certain threshold that marks the beginning of the tail region. In line with much discussed stylised features of financial returns, we find existence of tail-thickness in both the left and the right tails of the two innovation distributions, as indicated by positive values of the tail index of the

---

[1]For more on the methodology see NSE's website <www.nseindia.com>.

generalised Pareto distribution. Further, the left tail of each series is found to be heavier than the right tail, as indicated by higher values of the tail index for the left tail than that for the right tail. This clearly indicates tail asymmetry and negative skewness in the tails of both the series. These empirical findings provide interesting perspectives into our understanding of these markets, particularly in light of financial sector reforms initiated in India in 1990s that led to liberalisation and globalisation of financial markets.

The rest of this paper is organised as follows. Section 2 briefly discusses the Peaks-Over-Threshold (POT) model. Section 3 describes the data and methodology used in this paper. In Section 4 we present the empirical analysis of tail regions for Nifty (Subsection 4.1) and the IR series (Subsection 4.2) along with important results. Section 5 concludes this paper.

2   Peaks-Over-Threshold (POT) model

The POT model is based on "Pickands-Balkema-de Haan theorem" that postulates that the distribution of the observations in excess of certain high threshold can be approximated by a Generalised Pareto distribution (GPD). In this section we provide a brief description of the Pickands-Balkema-de Haan theorem and the Peaks-Over-Threshold model.

- **Pickands-Balkema-de Haan Theorem and GPD**

  Suppose that $X \in (l, u)$ is a random variable with density f and cdf F and k is a threshold on the range of the values of X. The Pickands-Balkema-de Haan theorem deals with the distribution of the excesses $(X - k)$ over certain high threshold k. It 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].

  The Generalised Pareto distribution (GPD) is given by

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

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

  The shape parameter ξ, called the tail index, determines the tail-thickness. When $\xi > 0$, we get the Fréchet distribution family, which is the class of distributions with regularly varying tails that include 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. When $\xi > 0$, moments of order $m > \frac{1}{\xi}$ are unbounded; thus $\frac{1}{\xi}$ determines the highest bounded moment for the distribution.

---

[2]Extremal-types theorem, also known as the Fisher-Tippett theorem is a fundamental result of extreme value theory. It identifies the possible limit laws for suitably normalised extreme observations from a sample of n i.i.d. realisations of a random variable X. Almost all common continuous distributions used in finance satisfy the extremal types theorem; hence the Pickands-Balkema-de Haan theorem is fairly general.

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 for which all moments exist. 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.

- **Tail estimation using Peaks-over-Threshold (POT) model**[3]

  In the POT model, first a threshold k is identified to define the start of the tail region. Then the distribution of the 'excesses' over the threshold point is estimated with the help of a GPD approximation.

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

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

  In terms of the underlying distribution F,

  $$\Phi_k(y) = \frac{F(y + k) - F(k)}{1 - F(k)} \tag{2}$$

  Applying Pickands-Balkema-de Haan theorem

  $$\Phi_k(y) \to G_{\xi, \beta(k)}(y) \quad \text{for} \quad k \to u \tag{3}$$

  Setting $x = k + y$ and using (2) and (3), we can rewrite F, for $x > k$, as

  $$F(x) = (1 - F(k))G_{\xi, \beta}(x - k) + F(k) \tag{4}$$

  Using empirical estimate for $F(k)$ and Maximum Likelihood (ML) estimates of the GPD parameters in (4) gives rise to the following tail estimator formula

  $$\hat{F}(x) = 1 - \frac{N_k}{N}\left(1 + \hat{\xi}\frac{x - k}{\hat{\beta}}\right)^{-\frac{1}{\hat{\xi}}} \tag{5}$$

  For a given probability level $p > F(k)$, a tail quantile is estimated by inverting the tail estimator formula (5),

  $$\hat{q}_p = k + \frac{\hat{\beta}}{\hat{\xi}}\left(\frac{N}{N_k}(1 - p)^{-\hat{\xi}} - 1\right) \tag{6}$$

  Where N_k is the number of observations beyond the threshold k.

  Equations (5) and (6) provide the basic formulae to estimate the tail probabilities and the tail quantiles.

---

[3]For a detailed treatment of POT model see standard textbooks on extreme value theory, eg. Embrechts et al. (1997).

3   Data and Methodology

In this section we describe the data and the methodology employed in the empirical analysis of the tail behaviour of two financial returns from India, viz., Nifty and 30-day IR.

The Nifty data consists of 2475 daily percentage logarithmic returns of the Nifty index from 7 November 1994 till 21 October 2004. The IR data consist of 1703 daily percentage logarithmic returns of the average of 30-day MIBID and MIBOR rates covering the period from 1 December 1998 till 3 November 2004. All data are obtained from the web site of NSE.

Table 1 presents some descriptive statistics (Panel A) and test statistics corresponding to tests of normality, skewness, kurtosis and autocorrelation (Panel B) for the two series. Normality of the data is tested by the Jarque-Bera test and the Anderson-Darling test. The parametric Jarque-Bera test is a joint test of the hypotheses of "zero skewness" and "zero excess kurtosis" (Jarque and Bera, 1980). The non parametric Anderson-Darling test is based on the empirical distribution function (EDF) and among all the well known tests based on EDF, Anderson-Darling test has the highest power in testing for normality against a wide range of alternatives when the parameters are unknown (Stephens, 1974). We have also carried out individual tests for "zero skewness", "zero excess kurtosis", and the Ljung-Box Portmanteau test (upto 35 lags) for the existence of autocorrelation[4].

The descriptive statistics clearly indicate heavy tails and negative skewness for the two returns series. These features are confirmed by the highly significant values of the test statistics corresponding to Jarque-Bera, Anderson-Darling, "zero skewness" and "zero kurtosis". As far as autocorrelation is concerned, Ljung-Box statistic is found to be significant for the Nifty returns but not significant for the IR returns.

The features of the Nifty and IR returns are graphically depicted in Figure 1. The top panel of Figure 1 depicts the QQ-plots of the two returns series against standard normal distribution. The QQ-plots indicate that the data are non-normal and their tails are heavy and asymmetric. The tail asymmetry is more profound in the case of the IR returns. The existence of serial correlation in the first and second moments of the returns series can be observed by looking at the correlograms of the returns and squared returns. In the middle and the bottom panels of Figure 1 we present correlograms for the returns and squared returns corresponding to Nifty (middle panel) and IR (bottom panel) series. As indicated by these correlograms, there is significant autocorrelation in the returns as well as the squared returns of Nifty series, indicating existence of time dynamics in the mean as well as volatility in the Nifty returns that can be represented by an ARMA-GARCH type model. For the IR series, there is no significant autocorrelation in the returns but the squared returns exhibit significant autocorrelation of order 1, thus indicating an ARCH(1)-type behaviour of the data.

Empirical methodology

Having described the essential features of the data, we describe below the various steps involved in the estimation of tail regions of the two series.

- *Time series specification*

---

[4]These standard tests can be found in any text book, eg., Campbell et al. (1997) (page 17 and page 47.)

It is important to have an appropriate 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.

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

$$r_t = \mu_t + \sigma_t z_t \tag{7}$$

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 with cumulative distribution function F.

The distribution of z_t, particularly its tail regions, is the focus of this paper. Application of POT model 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 white noise and does not contain any time dependence.

We use a pseudo maximum likelihood (PML) approach to estimate the parameters of mean and the volatility dynamics of the returns, using normal distribution for the innovation z_t. Under the PML method, use of normal distribution for the estimation does not imply the assumption of normality for the distribution of z_t. Under standard regularity conditions (Gourieroux, 1997; Gourieroux et al., 1984) the use of normal distribution would yield consistent estimates even if the underlying distribution is not normal.[5]

- *Choice of threshold*

  Having specified an appropriate time series model for each of the series, we extract the standard residuals coming out of the fitted model and use these residuals for estimating the tails of the innovation distribution.

  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. 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 beyond the threshold for estimating the GPD parameters, leading to statistical imprecision and very high variance of the estimates[6].

  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[7], 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 Neftci (2000) the threshold level is 1.645 times the unconditional

---

[5]Pseudo Maximum Likelihood (PML) estimators are obtained by maximising the likelihood function associated with a family of probability distributions that may not necessarily include the true distribution of the underlying random variable whose parameters are being estimated. Gourieroux et al. (1984) have established that the PML estimators of the first two moments (of the unknown distribution) based on the linear and quadratic exponential family are asymptotically consistent and normally distributed regardless of the exact form of the true unknown distribution. The normal distribution, being a quadratic exponential family, can provide consistent estimators of the first two moments. Moreover, this estimator is asymptotically normal.

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

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

variance of the data, which represents 5% of the extreme observations if the data were normally distributed.

In this paper, we follow a slightly different approach.[8] We first estimate the GPD parameters corresponding to various threshold levels, representing respectively 1%, 2%, 3%,.....till 15% of the extreme observations. Then we plot a graph of the estimated parameters and choose the threshold level at which the estimate stabilises. This is a non-parametric way of choosing the optimal threshold level and it is useful when the the mean-excess-plot or normal distribution assumption fail.

- *The Kolmogorov Smirnov test for discrepancy*

To test for the significance of difference between the estimated and empirical tails, we apply a non-parametric Kolmogorov-Smirnov (KS) test[9].

4 Empirical results

4.1 Nifty index

- *Time series specification for Nifty returns*

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.

Panel A of Table 2 presents the estimated parameters of the mean and volatility equations of the Nifty returns. All the parameters except for the constant in the mean equation, are found to be significantly different from zero. Thus, the Nifty returns have a significant AR(1) component in the mean dynamics and significant ARCH and GARCH effects of order 1 each, in the volatility dynamics, apart from a constant term in the volatility equation.

- *The standard residuals*

Having specified the time series dynamics of the data, we extract the standard residuals or the innovations z_t. This residual series is used for estimating the tail regions. Before the tail estimation, we carry out a preliminary statistical analysis on the series, in order to ascertain the statistical properties of the series.

Column 2 of Panel A of Table 3 presents summary statistics of the standard residuals obtained from the AR(1)-GARCH(1,1) specification of the Nifty returns. The values of these descriptive statistics indicate the existence of leptokurtosis and asymmetry in the standard residuals, although to a less extent compared to the returns series itself, as shown in Column 2 of Table 1. Thus, incorporation of conditional heteroskedasticity in the form of a GARCH(1,1) specification has only partially removed the leptokurtosis in the innovation distribution of Nifty returns. A substantial part of the leptokurtosis

---

[8]We tried with the mean excess plots but did not get a well behaved linear mean excess plot.
[9]See,eg. Hollander and Wolfe (1999).

still exists in the innovation distribution even after the time varying volatility has been removed.

Results of the tests of normality, skewness, kurtosis and autocorrelation on the standard residuals are presented in Column 2 of Panel B of Table 3. The test statistics corresponding to the tests of normality (Jarque-Bera and Anderson-Darling), skewness and kurtosis are found to be significantly different from zero at both 0.05 and 0.01 level of significance. However, the Ljung-Box statistic (upto lag 35) for test of autocorrelation is found to be insignificant at both 0.05 and 0.01 levels of significance. Thus, the innovation distribution can be regarded as asymmetric, leptokurtic and non-normal, *but* iid.

In Figure 2, we depict the features of the residuals by QQ-plot against the standard normal distribution and the correlograms of the residuals and squared residuals. A look at the QQ-plot of the Nifty residuals reveals the asymmetric and heavy tailed behaviour of the residuals. The correlograms in the middle panel of Figure 2 indicate that there is no statistical evidence of autocorrelation, both at the mean and at the volatility dynamics, in the residuals extracted from the Nifty series. Thus, the residual series of Nifty may be considered to be devoid of time series dynamics, and hence *iid*.

- *Choice of thresholds*

Defining the tail region involves the choice of an appropriate threshold level that marks the beginning of the tail. Choice of the threshold level is of critical importance so as to minimise the trade off between the bias and variance as discussed in Section 3.

In this paper, we first choose a number of thresholds and fit a Generalised Pareto distribution (GPD) to the excesses over each threshold. We use maximum likelihood method to estimate the parameters of the GPD. The final choice of the threshold is based on the graph of the estimated parameters as a function of the number of observations used for the estimation. The threshold corresponding to which the parameters stabilise is chosen as the appropriate threshold.

Figure 3 depict the plots of ML estimates of the GPD parameters ξ and β as a function of the number of extreme observations (N_k) used for the estimation. Panel A of this figure depict the estimates of ξ and β for the left tail and Panel B corresponds to the estimates of ξ and β for the right tail of the innovation distribution. As shown in the figures, the estimates are highly fluctuating when N_k is low, and reasonably stable as the the value of N_k increases.

In Panel A of Figure 3, the estimates of ξ and β behave in a stable manner from $N_k = 148$, which represents about 6% extreme negatives observations. Thus, for the estimation of the left tail, we choose a threshold level of 6% extreme negative observations.

Using Panel B of Figure 3 for choosing a threshold level for the right tail, we choose $N_k = 223$, the threshold level at which the estimates become stable. This represents 9% of the extreme positive observations.

- *Estimation of the tails*

Having chosen the threshold levels as described above, we present the estimates of the parameters of the GPD fitted to the excesses (peaks) over the respective thresholds in Panel A of Table 4. In this Table, we present, for each tail, the threshold points (k) corresponding to the chosen level, the number of extreme observations beyond the threshold (N_k) and the maximum likelihood estimates of the GPD parameters ξ and β along with their standard errors.

As shown in Panel A of Table 4, the estimated $\hat{\xi}$, the tail index is found to be positive for both the lower tail and the upper tail of the Nifty innovation distribution. This indicate that the tails on both sides of the innovation distribution are heavy. Further, the tail index estimated for the left tail (.2152) is higher than that for the right tail (.1188), implying that the left tail is heavier than the right tail. This tail asymmetry conforms to the rejection the test of skewness in Table 3.

The tail index determines the number of bounded moments of the distribution. This is given by the reciprocal of the higher tail index. In case of the Nifty innovation distribution, the number of bounded moments is the reciprocal of the left tail index; i.e.; $\frac{1}{0.2152} = 4.65$. This implies that the first four moments of the Nifty series are well defined, while moments of order higher than 4 are unbounded.

The estimates of the GPD parameters can be used in the tail quantile estimation formula (6) to estimate the tails of the distribution.

Table 5 presents some of the estimated quantiles, corresponding to fixed probability levels, on the left and the right tails of Nifty innovation, along with the empirical quantiles and the corresponding quantiles on the standard normal curve. 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 left tail and the Panel B presents the quantiles on the right tail.

As shown in Table 5, the estimated tail quantiles fit the empirical quantiles better than the normal distribution quantiles. This implies that a normal distribution approximation of the underlying DGP would provide to misleading estimation of extreme quantiles, and therefore a misleading interpretation of the risk.

Figure 5 presents a graph of the estimated tails compared with empirically observed tails and the corresponding normal approximation. As can be seen from this graphs, the estimated tails fits the empirical cdf better than the normal approximation. While normal approximation fails to capture the tail behaviour of the data on the upper tail, the generalised Pareto distribution is able to capture both the tails almost precisely.

In Panel A of Figure 7, we present the QQ-plot of the empirical tails against the generalised Pareto distribution for Nifty innovations. These plots indicate a very close fit of the respective tails by the generalised Pareto distribution.

- *The KS test for discrepancy*

  The close fit of the estimated tails to the empirical ones is statistically confirmed by

using a Kolmogorov-Smirnov (KS) test. Table 7 provides the estimated KS-statistics for testing the null hypothesis of no discrepancy between the estimated tail and the empirically observed tail. Panel A of this table refers to the Nifty 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.

4.2 The interest rates (IR) series

- *Time series model specification of the IR series*

  The correlograms of the IR series in the bottom panel of Figure 1 indicate that the series does not possess time series dynamics in the mean equation but there is evidence of an ARCH(1) type behaviour in the variance equation. Guided by this observation, we have estimated an ARCH(1) model for the returns of the IR series. The estimated parameters of the ARCH(1) model are presented in Panel B of Table 2. The ARCH(1) coefficient as well as the constant of the variance equation are found to be statistically significant.

- *The standard residuals*

  The standard residuals extracted from the estimated ARCH(1) model fitted to the IR series is found to be negatively skewed, leptokurtic, non-normal *but* independent of time dynamics. Descriptive statistics and test statistics for tests of normality, skewness, kurtosis and autocorrelation for the standard residuals are reported in column 3 of Panels A and B of Table 3. Similar pictures emerge from the QQ-plot (top panel of Figure 2, the right hand side graph) of the standard residuals and the correlograms (bottom panel of Figure 2) of the residuals and squared residuals. It is interesting to observe that the standard residuals have much higher coefficients of skewness and kurtosis than the raw returns of the IR series.

- *Choice of thresholds*

  As in the case of the Nifty series, we first fit generalised Pareto distribution to various threshold levels that represent 1%, 2%,....so on till 15% of the most extreme observations of the standard residuals. Then we draw the graphs of the estimated parameters as a function of the number of observations used for estimation. We choose the threshold level at which the estimates become stable.

  Figure 4 presents the graphs of the maximum likelihood estimates of parameters ξ and β as a function of the number of observations (N_k) used in estimating the generalised Pareto distribution to observations on left tail (Panel A) and right tail (Panel B) of the innovation distribution.

  Guided by these graphs, we choose $N_k = 102$ (representing 6% of the most negative observations) for the threshold level for the left tail and $N_k = 136$ (representing 8% of the most positive observations) for the right tail threshold level.

- *Estimation of the tails*

10

Panel B of Table 4 presents the parameter estimates of the GPD model fitted to the observations in excess (peaks) of the chosen thresholds. We report the threshold points on the tails (k), number of observations beyond threshold (N_k), the estimated cdf of the innovation at the chosen threshold (F_k) and the estimates of the GPD parameters.

According to these estimates, the left tail begins beyond probability level 0.06 of the distribution function and the right tail begins beyond probability level 0.92.

The positive value of the estimate of the tail index ξ for the left as well as the right tail indicate tail thickness for both the tails. In particular, the left tail with tail index value 0.3789 is much heavier than the right tail with tail index 0.0866. The highest number of bounded moments is given by the reciprocal of the tail index for the heavier tail (in this case the left tail) and it is 2.639. Thus, the innovation distribution of the IR returns has only the first two moments well defined, and the moments of order 3 and above are unbounded.

Using the estimates of the GPD parameters and equation (6) we have estimated some quantiles corresponding to fixed probability levels, on the left and right tails the innovation distribution of the IR returns. In Table 6 we present these estimated quantiles (in column 2), along with the empirically observed quantiles (Column 3) and the standard normal approximation (Column 4) corresponding to these probability levels, for a comparison. Panel A of this table presents tail quantiles on the left tail while Panel B presents quantiles on the right tail. A comparison of the estimated quantiles with the empirical and standard normal quantiles shows that the empirically observed quantiles are much closer to the estimated quantiles than the standard normal quantiles, particularly far in the tail, eg., at probability levels beyond 0.006 on the left tail and 0.994 on the right tail, where the normal distribution quantiles under-approximate the empirical quantiles.

Figure 6 depict the estimated tails of the innovation distribution of the IR returns, along with the empirical tails and the corresponding normal distribution tails. Panel A of this figure depict the left tail and Panel B depict the right tail. As shown by the graphs, the estimated tails almost precisely represent the empirically observed ones, while the normal distribution tails do not fit the empirical tails. Panel B of Figure 7 depict the QQ-plot of the empirical tails against the generalised Pareto distribution, which indicate the generalised Pareto fit to be a reasonable fit for the tails.

- *KS test*

  Panel B of Table 7 presents the KS test statistics for testing the discrepency between estimated and empirical tails. As shown by the non-significance of these statistics, there is no statistical evidence of discrepency.

4.3   Summarising these results

We summarise these empirical results as follows.

- The tails of both the innovations are asymmetric. The tail asymmetry for the two series are confirmed by the test of skewness as well as by different values of the tail indexes

for left and right tails of each series. Further, the tails of both series is found to be negatively skewed. This is indicated by the higher value of the tail index for the left tail than that for the right tail.

Such asymmetric tail behaviour is not uncommon for financial returns which are generally described by EGARCH-type volatility models. Particularly the negative asymmetry conforms to the observation that "negative shocks tend to be more persistent than the positive shocks", also known as "leverage effect" (Black, 1976).

- The tails of Nifty and IR innovations display heavy tailed behaviour. The hypothesis of zero excess kurtosis is rejected for both the series. Further, the positive tail indexes corresponding to left and right tails of each innovation series indicate that the innovation series belongs to the class of "regularly varying" tails, that define heavy tailed distribution as those characterised by failure of moments of some high order. In the case of the Nifty innovation, first four moments are well defined but moments of order 5 and higher are found to be unbounded. For the IR innovations, only the first two moments are bounded.

5 Conclusion

In this paper we present an empirical analysis of the tail behaviours of the innovation distributions of two prominent Indian financial series, viz., the Nifty index and the 30-day money market interest rate.

Using extreme value theory we estimate the tails of each of these series, by fitting Generalised Pareto distribution to the tails of the data. The positive value of the tail index, the shape parameter of the generalised Pareto distribution, indicate tail thickness in the tails of both the series. Further, the left tail of each series is found to be heavier than the right tail, indicating negative skewness. These findings conform with the stylised features of financial asset returns that are often discussed in the empirical literature.

In the wake of the remarkable growth of financial markets in India, initiated by financial sector reforms in 1990s, these empirical findings provides important insights into the understanding of these markets in India, particularly to the global investors,the risk managers and the academic community as a whole.

References

Black F, 1976. Studies on stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association Business and Economical Section* 177–181.

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

Campbell JY, Lo AW, Mackinlay AC, 1997. *The Econometrics of Financial Markets*. Princeton University Press.

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.

Gourieroux C, 1997. *ARCH Models and Financial Applications*. Springer.

Gourieroux C, Monfort A, Trognon A, 1984. Pseudo Maximum Likelihood Methods: Theory. *Econometrica* **52**:681–700.

Hollander M, Wolfe DA, 1999. *Nonparametric statistical methods*. John Wiley, New York, 2nd edn.

Jarque CM, Bera AK, 1980. Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters* **6**:255–259.

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.

Stephens MA, 1974. EDF Statistics for Goodness of Fit and Some Comparisons. *Journal of American Statistical Association* **69**:730–737.

Table 1 Description of Nifty and IR returns: Summary statistics (Panel A) and results of some statistical tests (Panel B)

This table presents descriptive statistics of the returns from the Nifty and the 30-day IR series and the results of some statistical tests on them.

The top panel of the table presents the descriptive statistics and the bottom panel presents results of two tests of normality (Jarque-Bera and Anderson-Darling), individual tests of "zero skewness" and "zero excess kurtosis" and the Ljung-Box Portmanteau test (upto lag 35). In the bottom panel $CV_{0.05}$ and $CV_{0.01}$ indicates the critical values of the relevant test statistic at 5% and 1% significane level respectively.

<div align="center">Panel A: Some descriptive statistics</div>

	Nifty	IR
Min	-13.054	12.662
Max	9.934	10.045
Mean	0.013	-0.035
SD	1.623	1.355
Skewness	-0.249	-0.893
Kurtosis	7.699	16.964

<div align="center">Panel B: Tests of Normality, Skewness, Kurtosis and Autocorrelation</div>

Test	Test statistic for		$CV_{0.05}$	$CV_{0.01}$
	Nifty	IR		
Jarque-Bera	2302.20*	14054.230*	5.99	9.21
Andersen-Darling	32413.308*	2378.371*	0.75	1.04
Skewness=0	-5.052^*	-15.040^*	1.96	2.58
Kurtosis=3	47.714*	117.593*	1.96	2.58
Portmanteau (till lag 35)	58.330*	29.718	49.802	57.342

* indicates significance at 0.01 level of significance

Table 2 Results of the estimation of time series models for Nifty and IR series

This table presents the estimated parameters, their SE's and the corresponding confidence intervals of the time series models fitted to the Nifty returns (Panel A) and the IR returns (Panel B). The appropriate time series for the Nifty series is an AR(1) − GARCH(1, 1) model and that for the IR series it is an ARCH(1) model.

Parameter	Estimates	SE	Confidence bounds
Panel A: Time series model for Nifty			
The mean equation:			
Constant	0.051	0.028	(-0.003, 0.106)
AR(1)	0.130	0.022	(0.087,0.174)
The variance equation:			
Constant	0.117	0.02	(0.068,0.167)
ARCH	0.129	0.015	(0.100, 0.159)
GARCH	0.833	0.018	(0.797, 0.868)
Panel B: Time series model for IR			
The variance equation:			
Constant	1.134	0.051	(1.034, 1.235)
ARCH	0.455	0.054	(0.349 0.561)

Table 3 Description of Standard residuals (SR): Summary statistics (Panel A) and results of some statistical tests (Panel B)

This table presents descriptive statistics and the results of some statistical tests on the standard residuals (SR) extracted from an AR(1) − GARCH(1, 1) specification of the Nifty returns and an ARCH(1) specification of the IR returns.

The top panel of the table presents the descriptive statistics and the bottom panel presents results of two tests of normality (Jarque-Bera and Anderson-Darling), the individual tests of "zero skewness" and "zero excess kurtosis" and the Ljung-Box Portmanteau test (upto lag 35). In the bottom panel $CV_{0.05}$ and $CV_{0.01}$ indicates the critical values of the relevant test statistic at 5% and 1% significane level respectively.

Panel A: Some descriptive statistics

	Nifty	IR
Min	-6.312	-11.828
Max	6.576	4.860
Mean	-0.034	0.010
SD	0.999	1.001
Skewness	-0.128	-1.516
Kurtosis	5.984	22.479

Panel B: Tests of Normality, Skewness, Kurtosis and Autocorrelation

Test	Test statistic for		$CV_{0.05}$	$CV_{0.01}$
	Nifty SR	IR SR		
Jarque-Bera	924.714*	27527.030*	5.99	9.21
Andersen-Darling	2397.564*	1626.693*	0.75	1.04
Skewness=0	-2.598^*	-25.516^*	1.96	2.58
Kurtosis=3	30.298*	163.939*	1.96	2.58
Portmanteau (till lag 35)	48.060	35.393	49.802	57.342

* indicates significance at 0.01 level of significance

Table 4 Results of the GPD estimation

This table provides the results of the estimated GPD parameters fitted to the excesses over the chosen thresholds of the standard residuals of Nifty returns (Panel A) and of interest rate returns (Panel B).
The first column gives the threshold points on left and right tails. 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 Maximum-likelihood estimation of the GPD parameters fitted to the excesses over the thresholds, along with the standard errors of estimation within parenthesis.

	k	N_k	F_u	$\hat{\xi}$	$\hat{\beta}$
Panel A: Nifty innovation					
left tail	- 1.4885	148	0.06	0.2152	0.4826
				(0.0942)	(0.0597)
Right tail	1.2301	223	0.901	0.1188	0.4817
				(0.0698)	(0.0464)
Panel B: IR innovation					
left tail	-1.2734	102	0.06	0.3789	0.5739
				(0.1293)	(0.0909)
Right tail	1.1524	136	0.92	0.0866	0.6793
				(0.1100)	(0.0946)

Figures in parenthesis indicate standard error

Table 5 Tail quantiles (Nifty innovation) : estimated, empirical and standard normal

This table provides some of the estimated quantiles on the innovation distribution of Nifty returns for pre-specified probability levels, along with the empirical quantiles as well as the corresponding quantiles on the standard normal distribution. Panel A deals with the left tail and Panel B deals with the right tail. 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	Estimated	Empirical	Normal
Panel A: Quantiles on the left tail			
0.050	-1.5767643	-1.5842934	-1.6448536
0.040	-1.6914221	-1.7388918	-1.7506861
0.030	-1.8476038	-1.9022421	-1.8807936
0.020	-2.0848123	-2.068741	-2.0537489
0.010	-2.5414851	-2.3643709	-2.3263479
0.009	-2.6170602	-2.4016984	-2.3656181
0.008	-2.7035996	-2.484580	-2.4089155
0.007	-2.8043997	-2.5759851	-2.4572634
0.006	-2.9244247	-2.8300880	-2.5121443
0.005	-3.0716206	-3.2417906	-2.5758293
0.004	-3.2598135	-3.3000028	-2.6520698
0.003	-3.5161617	-3.5265452	-2.7477814
0.002	-3.9055028	-4.0759153	-2.8781617
0.001	-4.6550607	-5.1706948	-3.0902323
Panel B: Quantiles on the right tail			
0.950	1.5241429	1.5173867	1.6448536
0.960	1.6409678	1.6347701	1.7506861
0.970	1.7962243	1.7906302	1.8807936
0.980	2.0242547	1.9504204	2.0537489
0.990	2.4404396	2.4621405	2.3263479
0.991	2.5067556	2.5714443	2.3656181
0.992	2.5818798	2.6267498	2.4089155
0.993	2.6683298	2.7514615	2.4572634
0.994	2.7698494	2.7915322	2.5121443
0.995	2.8923463	2.8712860	2.5758293
0.996	3.0459265	3.0623585	2.6520698
0.997	3.2500297	3.4957898	2.7477814
0.998	3.5498028	3.6842619	2.8781617
0.999	4.0969276	4.4270124	3.0902323

Table 6 Tail quantiles (IR innovation) : estimated, empirical and standard normal

This table provides some of the estimated quantiles on the innovation distribution of the interest rate returns for pre-specified probability levels, along with the empirical quantiles as well as the corresponding quantiles on the standard normal distribution. Panel A deals with the left tail and Panel B deals with the right tail. 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	Estimated	Empirical	Normal
Panel A: Quantiles on the left tail			
0.050	-1.3816945	-1.3971946	-1.6448536
0.040	-1.3941658	-1.4372644	-1.7506861
0.030	-1.7282857	-1.6887784	-1.8807936
0.020	-2.0553533	-2.0127567	-2.0537489
0.010	-2.7451491	-2.7242100	-2.3263479
0.009	-2.8667826	-2.8998683	-2.3656181
0.008	-3.0086311	-2.9247548	-2.4089155
0.007	-3.1772915	-2.9457809	-2.4572634
0.006	-3.3829092	-3.4370357	-2.5121443
0.005	-3.6421250	-3.4415974	-2.5758293
0.004	-3.9847438	-4.0022029	-2.6520698
0.003	-4.4714363	-4.3958650	-2.7477814
0.002	-5.2540321	-4.9042673	-2.8781617
0.001	-6.9045510	-7.0330868	-3.0902323
Panel B: Quantiles on the right tail			
0.950	1.4783245	1.4532740	1.6448536
0.960	1.6377382	1.5856886	1.7506861
0.970	1.8478569	1.8424819	1.8807936
0.980	2.1530334	2.1085473	2.0537489
0.990	2.7002072	2.9705064	2.3263479
0.991	2.7862927	3.1767720	2.3656181
0.992	2.8834625	3.2037110	2.4089155
0.993	2.9948300	3.2416549	2.4572634
0.994	3.1250060	3.3244496	2.5121443
0.995	3.2812317	3.4651166	2.5758293
0.996	3.4758240	3.7049252	2.6520698
0.997	3.7323106	3.8585384	2.7477814
0.998	4.1048319	4.4807940	2.8781617
0.999	4.7727531	4.5060569	3.0902323

Table 7 Results of the Kolmogorov-Smirnov (KS) tests

This table provides the values of the Kolmogorov-Smirnov test statistics for testing the hypotheses of "no discrepency" between the estimated and the empirical tails of the innovations of Nifty (Panel A) and IR series (Panel B). Values of the KS statistic for each tail in column 2 are reported along with the critical values at 5% and 1% levels of significance in columns 3 and 4 respectively. The critical values of KS statistic depend on the sample size.

	KS statistic	$CV_{0.05}$	$CV_{0.01}$
Panel A: KS test for Nifty innovations			
Left tail	0.0038	0.1118	0.1339
Right tail	0.0042	0.0911	0.1092
Panel B: KS test for IR innovations			
Left tail	0.00366	0.13466	0.161394
Right tail	0.00362	0.11661	0.13977

Figure 1 QQ-plots and Correlograms of Nifty and IR returns

These figures graphically describe the features of Nifty and IR returns. The top panel presents the QQ-plots against standard normal distribution of the two series. The middle and the bottom panels depict correlograms of the returns and squared returns of the two series (middle panel for Nifty and bottom panel for IR).
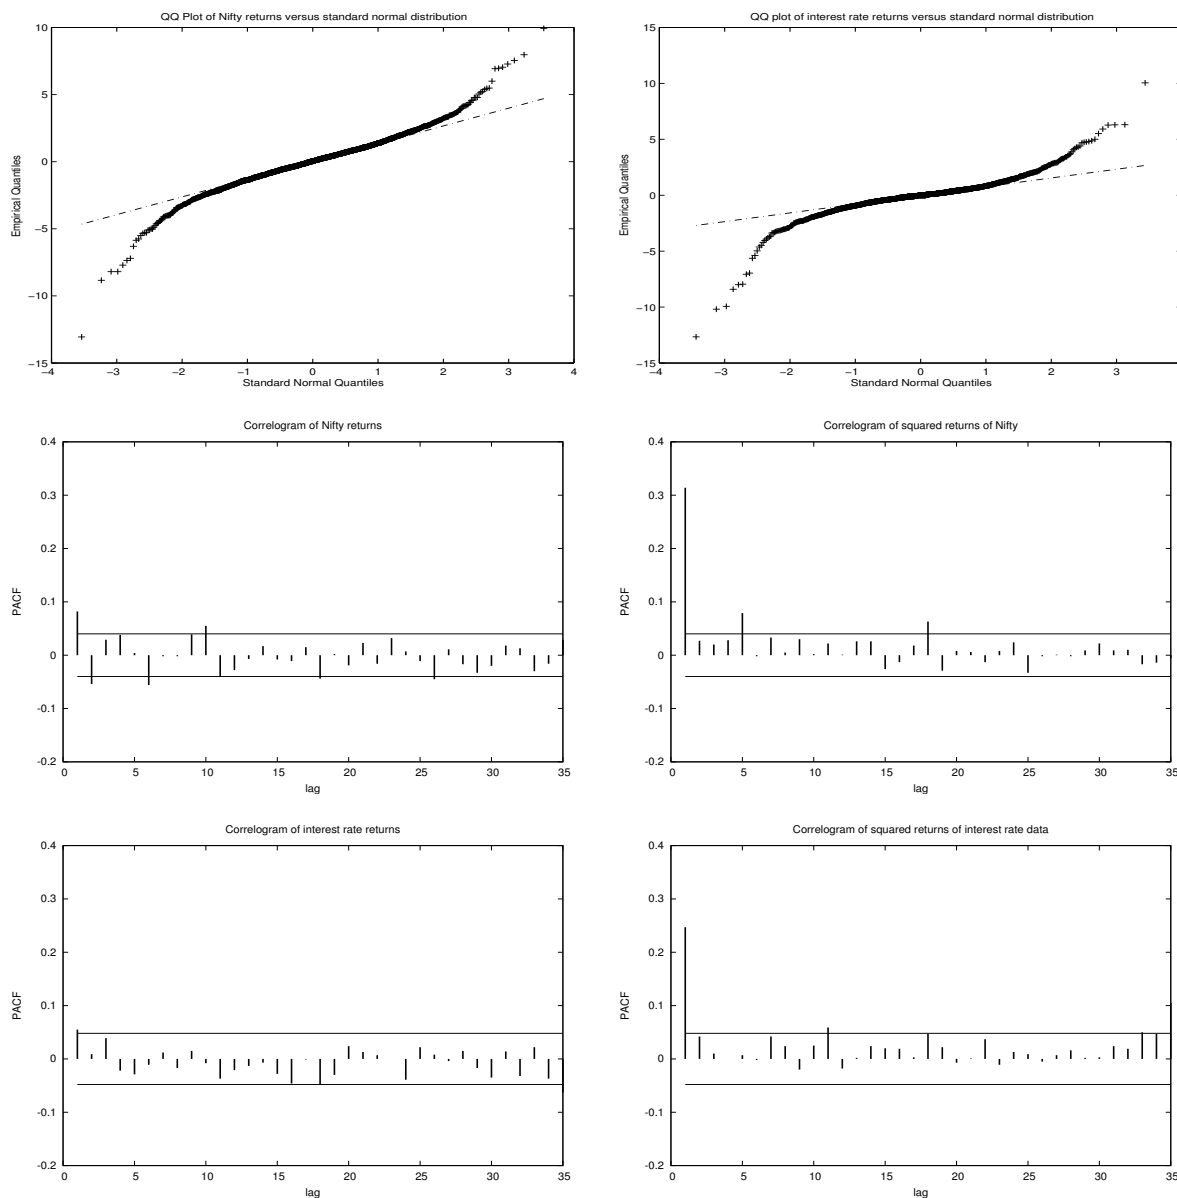
Figure 2 QQ-plots and Correlograms of the standard residuals of Nifty and IR series

These figures graphically describe the features of the standard residuals (SR) obtained from the respective time series model fitted to the Nifty and IR returns. The top panel presents the QQ-plots against standard normal distribution of the two residual series. The middle and the bottom panels depict correlograms of the residuals and squared residuals of the two series (middle panel for Nifty residuals and bottom panel for IR residuals).
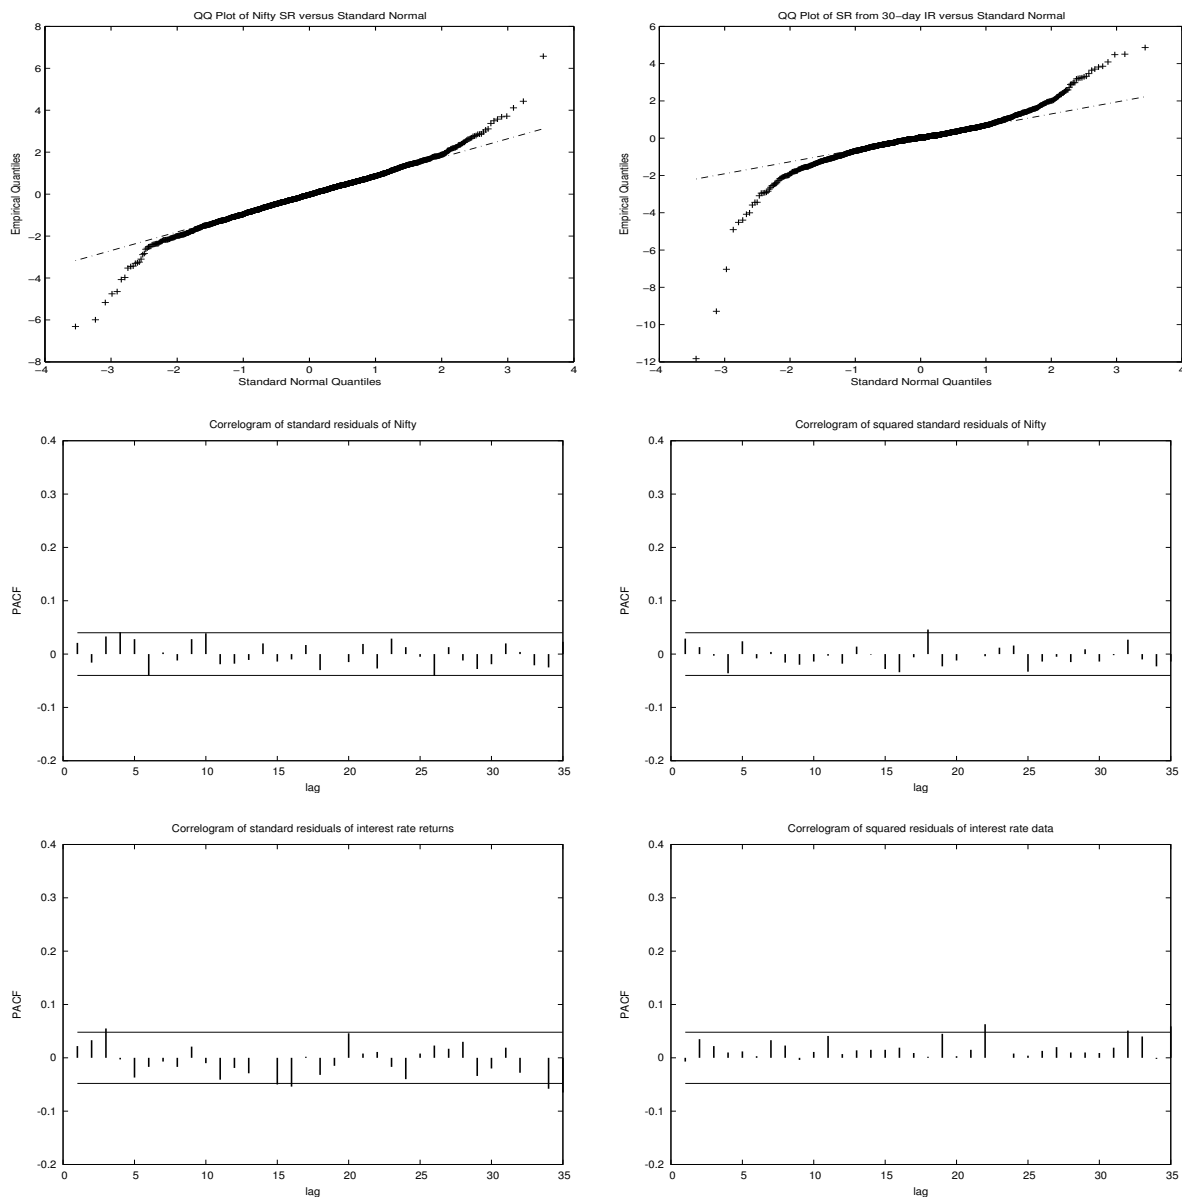
Figure 3 Plots of $\hat{\xi}$ and $\hat{\beta}$ against various threshold levels: Nifty data

These figures depict the maximum likelihood estimates of the GPD parameters $\hat{\xi}$ and $\hat{\beta}$ corresponding to various threshold levels on the tails of the innovation distribution of Nifty returns. The top panel corresponds to the estimates for various threshold levels on the left tail and the bottom panel corresponds to the estimates for the right tail.
In each of these graphs, the solid line represents the estimated value of the parameters while the dotted line represent the confidence bounds.

Panel A: Estimates of ξ and β for various thresholds on the left tail

Panel B: Estimates of ξ and β for various thresholds on the right tail

This figures depict the maximum likelihood estimates of the GPD parameters $\hat{\xi}$ and $\hat{\beta}$ corresponding to various threshold levels on the tails of the innovation distribution of the interest rate returns. The top panel corresponds to the estimates for various threshold levels on the left tail and the bottom panel corresponds to the estimates for the right tail.

In each of these graphs, the solid line represents the estimated value of the parameters while the dotted line represent the confidence bounds.

Panel A: Estimates of ξ and β for various thresholds on the left tail

Panel B: Estimates of ξ and β for various thresholds on the right tail
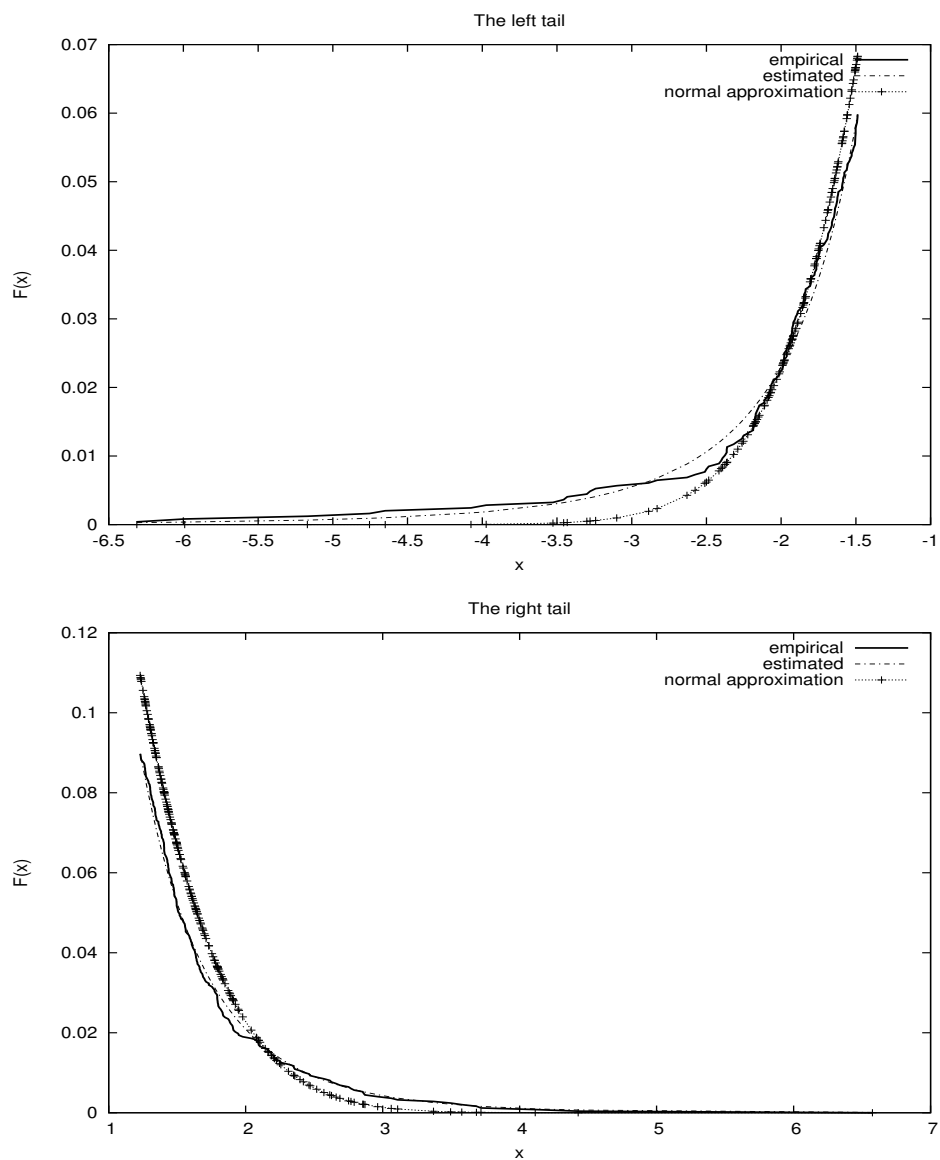
Figure 5 Tails of the innovation distribution of Nifty returns

These figures depict the left tail (the graph on the top) and the right tail (the graph on the bottom) of the innovation distribution of Nifty returns. In both the graphs, we draw estimated cdf, empirical cdf and the normal approximations corresponding to the observed data.
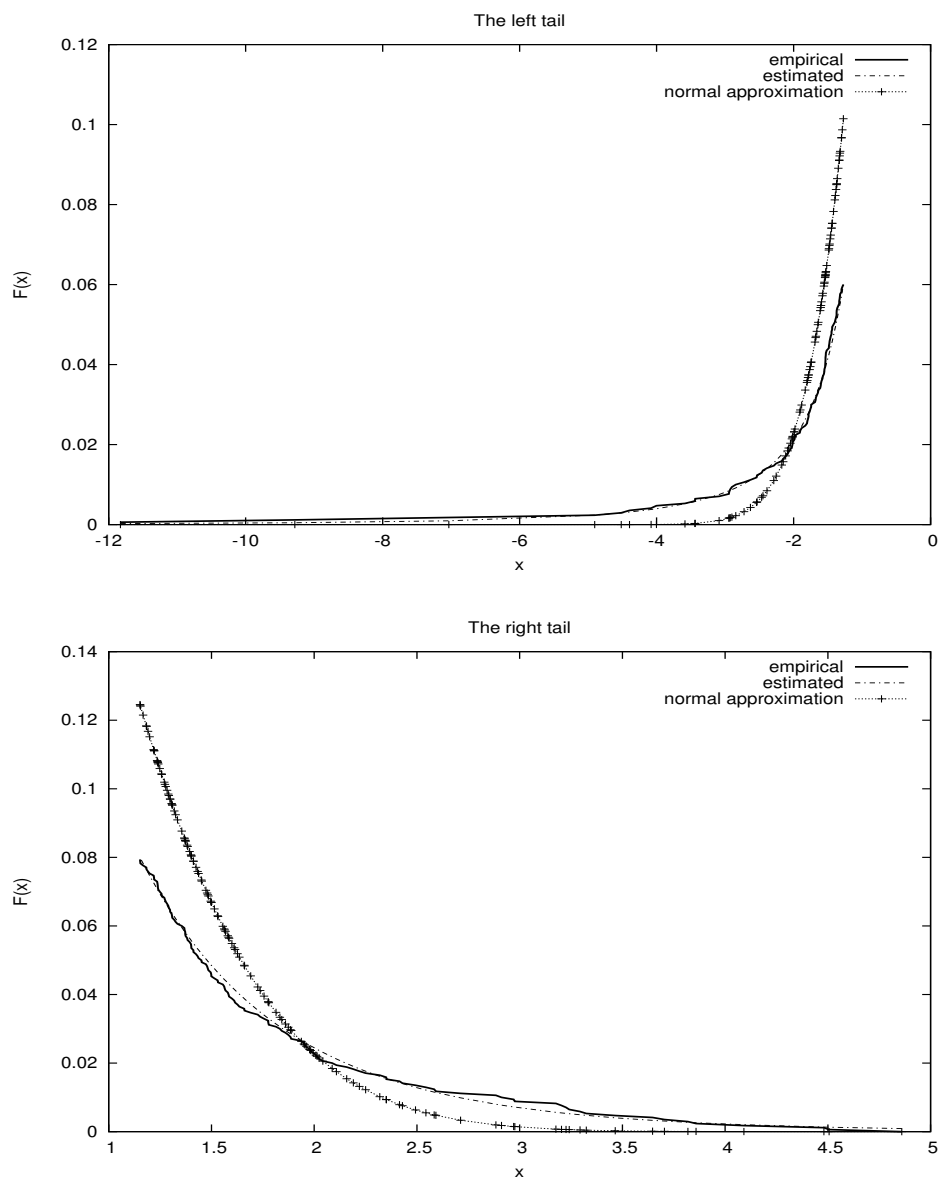
Figure 6 Tails of the innovation distribution of Interest rate returns

These figures depict the left tail (the graph on the top) and the right tail (the graph on the bottom) of the innovation distribution of interest rate returns. In both the graphs, we draw estimated cdf, empirical cdf and the normal approximations corresponding to the observed data.

Figure 7 QQ-plots of the tails against Generalised Pareto Distribution

This figures depict the QQ-plot of the tails of the Nifty innovation against the generalised Pareto distribution (the graphs on the top panels) and that of the tails of the interest rate innovations (the graphs on the bottom).

Panel A: QQ-plot of tails for Nifty data

Panel B: QQ-plot of tails for 30-day IR data