1 An Empirical Examination of Daily Stock Return Distributions for U.S. Stocks

Svetlozar T. Rachev*[1], Stoyan V. Stoyanov[2], Almira Biglova[3], and Frank J. Fabozzi[4]

[1] Department of Econometrics and Statistics, University of Karlsruhe, D-76128 Karlsruhe, Germany and Department of Statistics and Applied Probability, University of California Santa Barbara, CA 93106, USA
[2] Chief Financial Researcher (FinAnalytica, Inc., Seattle, USA)
[3] Researcher, Institute of Econometrics, Statistics and Mathematical Finance School of Economics and Business Engineering University of Karlsruhe
[4] Frederick Frank Adjunct Professor of Finance, Yale University, School of Management

Abstract. This article investigates whether the Gaussian distribution hypothesis holds 382 U.S. stocks and compares it to the stable Paretian hypothesis. The daily returns are examined in the framework of two probability models - the homoskedastic independent, identical distributed model and the conditional heteroskedastic ARMA-GARCH model. Consistent with other studies, we strongly reject the Gaussian hypothesis for both models. We find out that the stable Paretian hypothesis better explains the tails and the central part of the return distribution.

Keywords Stable distributions, ARMA-GARCH, assets returns, heavy tails

1.1 Introduction

The cornerstone theories in finance such as mean-variance model for portfolio selection and asset pricing models that have been developed rest upon the assumption that asset returns follow a normal distribution. Yet, there is little, if any, credible empirical evidence that supports this assumption for financial assets traded in most markets throughout the world. Moreover, the evidence is clear that financial return series are heavy-tailed and, possibly, skewed. Fortunately, several papers have analyzed the consequences of relaxing the normality assumption and developed generalizations of prevalent concepts in financial theory that can accommodate heavy-tailed returns (see Rachev and Mittnik, 2000 and Rachev (2003) and references therein).

Mandelbrot (1963) strongly rejected normality as a distributional model for asset returns, conjecturing that financial return processes behave like

* Prof Rachev gratefully acknowledges research support by grants from Division of Mathematical, Life and Physical Sciences, College of Letters and Science, University of California, Santa Barbara, the Deutschen Forschungsgemeinschaft and the Deutscher Akademischer Austausch Dienst.

non-Gaussian stable processes. To distinguish between Gaussian and non-Gaussian stable distributions. The latter are commonly referred to as "stable Paretian" distributions or "Levy stable" distributions [1].

While there have been several studies in the 1960s that have extended Mandelbroit's investigation of financial return processes, probably, the most notable is Fama (1963, 1965). His work and others led to a consolidation of the stable Paretian hypothesis. In the 1970s, however, closer empirical scrutiny of the "stability" of fitted stable Paretian distributions also produced evidence that was not consistent with the stable Paretian hypothesis. Specifically, it was often reported that fitted characteristic exponents (or tail-indices) did not remain constant under temporal aggregation[2]. Partly in response to these empirical "inconsistencies", various alternatives to the stable law were proposed in the literature, including fat-tailed distributions being only in the domain of attraction of a stable Paretian law, finite mixtures of normal distributions, the Student t-distribution, and the hyperbolic distribution.

A major drawback of all these alternative models is their lack of stability. As has been stressed by Mandelbrot and argued by Rachev and Mittnik (2000), among others, the stability property is highly desirable for asset returns. This is particularly evident in the context of portfolio analysis and risk management. Only for stable distributed returns does one obtain the property that linear combinations of different return series (e.g., portfolios) follow again a stable distribution. Indeed, the Gaussian law shares this feature, but it is only one particular member of a large and flexible class of distributions, which also allows for skewness and heavy-tailedness.

Recent attacks on Mandelbrot's stable Paretian hypothesis focus on the claim that empirical asset return distributions are not as heavy-tailed as the non-Gaussian stable law suggests. Studies that come to such conclusions are typically based on tail-index estimates obtained with the Hill estimator. Because sample sizes beyond 100,000 are required to obtain reasonably accurate estimates, the Hill estimator is highly unreliable for testing the stable hypothesis. More importantly, the Mandelbrot's stable Paretian hypothesis is interpreted too narrowly, if one focuses solely on the *marginal* distribution of return processes. The hypothesis involves more than simply fitting marginal asset return distributions. Stable Paretian laws describe the fundamental "building blocks" (e.g., innovations) that drive asset return processes. In addition to describing these "building blocks," a complete model should be rich enough to encompass relevant stylized facts, such as (1) non-Gaussian, heavy-tailed and

[1] Stable Paretian is used to emphasize that the tails of the non-Gaussian stable density have Pareto power-type decay "Levy stable" is used in recognition of the seminal work of Paul Levy's introduction and characterization of the class of non-Gaussian stable laws.

[2] For a more recent study, see Akgiray and Booth (1988) and Akgiray, Lamoureux (1989).

skewed distributions, (2) volatility clustering (ARCH-effects), (3) temporal dependence of the tail behavior, and (4) short- and long-range dependence.

An attractive feature of stable models — not shared by other distributional models — is that they allow us to generalize Gaussian-based financial theories and, thus, to build a coherent and more general framework for financial modeling. The generalizations are only possible because of specific probabilistic properties that are unique to (Gaussian and non-Gaussian) stable laws, namely, the stability property, the Central Limit Theorem and the Invariance Principle for stable processes [3].

In this paper we present additional empirical evidence comparing normal and stable Paretian for a large sample of U.S. stocks. Our empirical analyses go beyond those typically found in the literature wherein the focus is almost exclusively on the *unconditional* distribution of equity returns. Here we also investigate conditional homoskedastic (i.e., constant-conditional-volatility) and heteroskedastic (i.e., varying-conditional-volatility) distributions. It is because asset returns typically exhibit temporal dependence that the *conditional* distributions are of interest. If asset return embed information on past market movements, it is not the *unconditional* return distribution which is of interest, but the *conditional* distribution, which is conditioned on information contained in past return data, or a more general information set.

We believe our study is the first one investigating the stable Paretian distribution for equity returns that includes a large sample of companies (382). Most other studies have been limited to stock indexes. In studies where individual stock returns have been analyzed, the samples have been small, typically limited to the constituent components of the Dow Jones Industrial Average or a non-U.S. stock index with no more than 40 stocks. Most likely the reason for the limitation of other researchers to use a large sample of companies to analyze the stable Paretian distribution is the computational time involved to calculate the maximum likelihood estimate of the parameters of the stable distribution. For our 382 companies, for example, it took approximately 3.5 hours on PC 2.4 GHz Intel Pentium IV, 1Gb of RAM to compute the stable parameters for both models that we estimate in this study. All calculations were done in MATLAB. The same amount of calculations would have taken about 17.5 hours in 1999 and about 68 hours in 1996. The RAM requirements would have made such a study very hard to organize in 1990 and practically impossible in 1985 on an average PC [4].

[3] Detailed accounts of properties of stable distributed random variables can be found in Samorodnitsky and Taqqu (1994) and Janicki and Weron (1994).

[4] DuMouchel (1971) was the first to study maximum likelihood estimation of the parameters of stable distributions in his dissertation in the beginning of the 1970s. He did his calculations on a CDC 6400 computer at the University of California, Berkeley. The computers of this class were among the fastest computers in the world at the end of the 1960s. Assuming infinite amount of computer memory available, he would have needed more than 20 days to perform only the parameter fitting of the stable distributions for our sample. This extremely rough estimate

The paper is organized as follows. The methodology employed is explained in Section 1.2 followed by a description of our sample in Section 1.3. The empirical results are reported in Section 1.4 and a summary of our major conclusions are presented in Section 1.5.

1.2 Methodology

As noted in Section 1.1, because asset returns typically exhibit temporal dependence, the focus of the analysis should be on *conditional* distributions. The class of autoregressive moving average (ARMA) models is a natural candidate for conditioning on the past of a return series. These models have the property that the conditional distribution is homoskedastic. However, because financial markets frequently exhibit volatility clusters, the homoskedasticity assumption may be too restrictive. As a consequence, conditional heteroskedastic models, such as that proposed by autoregressive conditional heteroskedastic (ARCH) models as proposed Engle (1982) and the generalized GARCH proposed by Bollerslev (1986), possibly in combination with an ARMA model, referred to as an ARMA-GARCH model, are common in empirical finance. It turns out that ARCH-type models driven by normally distributed innovations imply unconditional distributions which themselves possess heavier tails. However, many studies have shown that GARCH-filtered residuals are themselves heavy-tailed, so that stable Paretian distributed innovations ("building blocks") would be a reasonable distributed assumption.

In the general case, no closed-form expressions are known for the probability density and distribution functions of stable distributions. They are described by four parameters: α, called the index of stability, which determines the tail weight or density's kurtosis with $0 < \alpha \leq 2$, β, called the skewness parameter, which determines the density's skewness with $-1 \leq \beta \leq 1$, $\sigma > 0$ which is a scale parameter, and μ which is a location parameter. Stable distributions allow for skewed distributions when $\beta \neq 0$ and when β is zero, the distribution is symmetric around μ. Stable Paretian laws have fat tails meaning that extreme events have high probability relative to the normal distribution when $\alpha < 2$. The Gaussian distribution is a stable distribution with $\alpha = 2$. (For more details on the properties of stable distributions see Samorodnitsky, Taqqu (1994).) Of the four parameters, α and β are most important as they identify two fundamental properties that are untypical of the normal distribution- heavy tails and asymmetry.

1.3 Description of the data

The sample of companies used in this study was obtained as follows. We began with all the companies that were included in the S&P500 index over the 12-

---

was determined by us based on the calculation time in DuMouchel's dissertation and the relative speed of his algorithm compared to ours.

year time period January 1, 1992 to December 12, 2003. The constituent companies in the S&P 500 are determined by a selection committee of the Standard & Poor's Corporation which periodically adds and removes companies from the index. Over the 12-year time period, there were more than 800 companies that had been included in the index. We then selected from these companies those that had a complete return history (3,026 observations). The rest of the companies that were included in S&P500 stocks but not in our sample have shorter historical series with unequal histories.

Daily returns were calculated as $r(t) = \log(S(t)/S(t-1))$, where $S(t)$ is the stock value at t (the stocks are adjusted for dividends).

1.4 Tests and results

We employ two tests in our investigation. In the first test, we assume that daily return observations are independent and identically distributed (iid); in the second test, the daily return observations are assumed to follow a ARMA(1,1)-GARCH(1,1) process. The first test concerns the unconditional, homoskedastic distribution model while the second one belongs to the class of conditional heteroskedastic models.

For both tests, we verify whether the Gaussian hypothesis holds. The normality tests employed are based on the Kolmogorov distance (KD) and computed according to

$$KD = \sup_{x \in \mathbb{R}} |F_s(x) - F(x)|$$

where $F_s(x)$ is the empirical sample distribution and $F(x)$ is the cumulative distribution function (cdf) of the estimated parametric density and emphasizes the deviations around the median of the distribution.

For both the iid and the ARMA-GARCH tests, we compare the goodness-of-fit in the case of the Gaussian hypothesis and the more general stable Paretian hypothesis. We use two goodness-of-fit measures for this purpose, the KD-statistic and the Anderson-Darling (AD) statistic. The AD-statistic is computed as follows:

$$AD = \sup_{x \in \mathbb{R}} |F_s(x) - F(x)|$$

The AD-statistic accentuates the discrepancies in the tails. Since in the calculation of the AD statistic the extreme observations are most important, we repeat the calculations assuming that the observations below the 0.1% quantile and above 99.9% quantile result from errors in the data. We replace them with the average of the two adjacent observations.

1.4.1 The iid model

In the simple setting of the iid (independent identically distributed returns) model, we have estimated the values for the four parameters of the stable Paretian distribution using the method of maximum likelihood. Figure 1.1 shows a scatter plot of the estimated pairs (α, β) for all stocks.

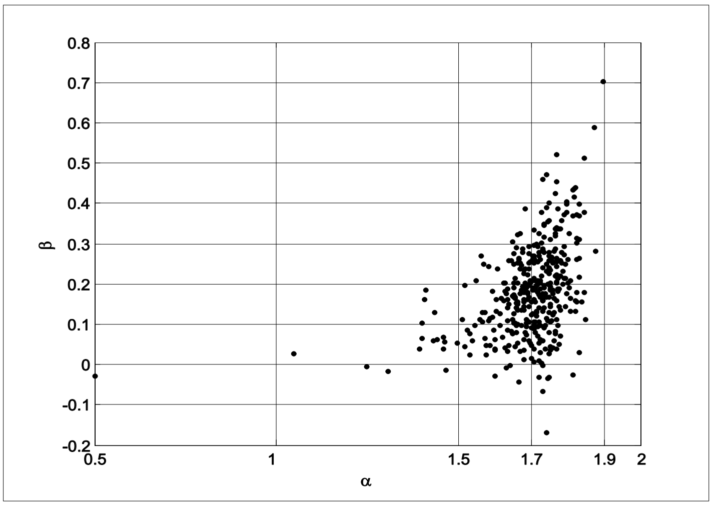

Fig. 1.1. Scatter plot of the index of stability and the skewness parameter for the daily returns of 382 stocks, iid model.

We find that for the return distribution of every stock in our sample that (1) the estimated values of the index of stability are below 2 and (2) there is asymmetry ($\beta \neq 0$). These two facts alone would strongly suggest that the Gaussian assumption is not a proper theoretical distribution model for describing the return distribution of stocks. Additional support for the stable Paretian hypothesis is contained in Tables 1.1 and 1.2. The two tables show that we can reject the normality using the standard Kolmogorov-Smirnov test for (1) more than 95% of the stocks at the extremely high confidence level of 99.9% and (2) 100% of the stocks at the traditional levels of 95% and 99%. In contrast, the stable-Paretian hypothesis is rejected in much fewer cases.

The superiority of the stable Paretian assumption over the Gaussian assumption is clearly seen by examining Figures 4 and 5. The figures show the computed KD and AD statistics for all stocks under the two distributional assumptions. For every stock in our sample, the KD-statistic in the stable Paretian case is below the KD-statistic in the Gaussian case. The same is true for the AD-statistic. The KD-statistic implies that for our sample firms there

Confidence level	95%	99%	99.50%	99.90%
Original data	100.00%	100.00%	99.74%	99.21%
Truncated data	100.00%	99.48%	98.43%	96.34%

Table 1.1. Percentage of stocks for which normality is rejected at different confidence levels using Kolmogorov distance in the iid model.

a better fit of the stable Paretian model around the center of the distribution while the AD-statistic implies a better fit in the tails. The huge difference between the AD-statistic computed for the stable Paretian model relative to the Gaussian model strongly suggests a much better ability for the stable Paretian model to forecast extreme events and confirms an already noticed phenomenon — the Gaussian distribution fails to describe observed large downward or upward asset price shifts, that is in reality extreme events have larger probability than predicted by the Gaussian distribution.

Confidence level	95%	99%	99.50%	99.90%
Original data	28.54%	15.71%	11.78%	7.07%
Truncated data	33.51%	16.49%	12.83%	7.85%

Table 1.2. Percentage of stocks for which the Stable Paretian hypothesis is rejected at different confidence levels using Kolmogorov distance in the iid model.

Summary statistics of the various statistical tests and parameter estimates for the entire sample are provided in Table 1.3. Truncating the data improves the AD statistic for the Gaussian model but it remains more than 8 times larger than the AD statistic for the stable Paretian model. In the case of the non-truncated data, the former is more than 10 times larger than the latter.

		α	β	KD Normal	KD Stable	AD Normal	AD Stable
	mean	1.697	0.18	0.0718	0.0241	0.6575	0.0612
Original	median	1.715	0.18	0.0638	0.0217	0.5867	0.0577
Data	25% quantile	1.661	0.1	0.0543	0.0184	0.5128	0.0517
	75% quantile	1.761	0.25	0.0769	0.0257	0.8084	0.0639
	mean	1.72	0.2	0.0588	0.0248	0.5977	0.0669
Truncated	median	1.739	0.19	0.0543	0.0223	0.5128	0.064
Data	25% quantile	1.683	0.11	0.0471	0.0193	0.365	0.0581
	75% quantile	1.786	0.27	0.0633	0.0268	0.8084	0.0696

Table 1.3. Summary statistics for the entire sample of 382 stocks, iid model.

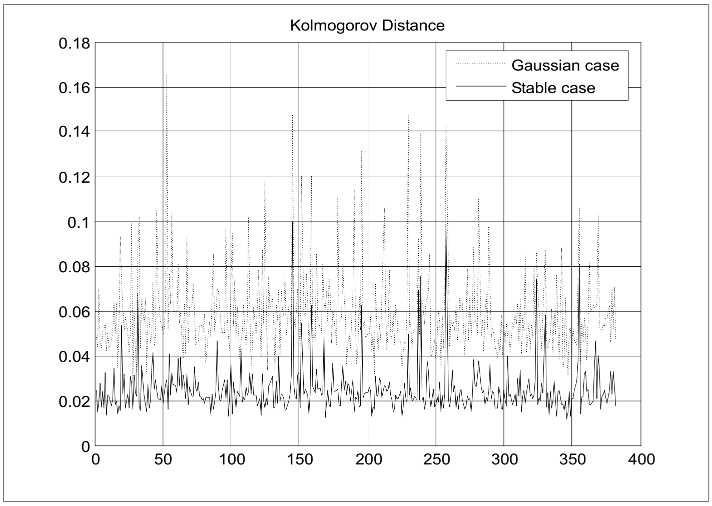

Fig. 1.2. Kolmogorov Distance in both stable Paretian and Gaussian cases for all 382 stocks, iid model.

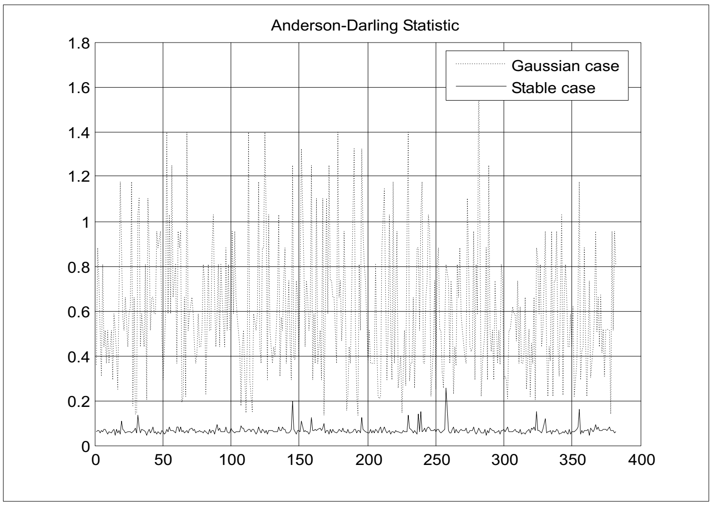

Fig. 1.3. Anderson-Darling statistic in both stable Paretian and Gaussian cases for all 382 stocks, iid model.

1.4.2 ARMA-GARCH model

Since the simple iid model does not account for the clustering of the volatility phenomenon, as a second test, we consider the more advanced ARMA-GARCH model. The general form of the ARMA(p,q)-GARCH(r,s) model is:

$$\begin{aligned} r_t &= C + \sum_{i=1}^{p} a_i r_{t-i} + \sum_{i=1}^{q} b_i \epsilon_{t-i} + \epsilon_t \\ \epsilon_t &= \sigma_t \delta_t \\ \sigma_t^2 &= K + \sum_{i=1}^{r} w_i \epsilon_{t-i}^2 + \sum_{i=1}^{s} \nu_i \sigma_{t-i}^2 \end{aligned} \tag{1.1}$$

where a_i, $i = 1, \ldots, p$, b_i, $i = 1, \ldots, q$, w_i, $i = 1, \ldots, r$, ν_i, $i = 1, \ldots, s$, C and K are the model parameters. δ_t's are called the *innovations process* and are assumed to be iid random variables which we additionally assume to be either Gaussianl or stable Paretian. An attractive property of the ARMA-GARCH process is that it allows a time-varying volatility via the last equation.

The model in which $p = q = r = s = 1$ proved appropriate for the 382 stock returns time series that we consider because the serial correlation in the residuals disappeared. The model parameters are estimated using the method of maximum likelihood assuming the normal distribution for the innovations. In this way, we maintain strongly consistent estimators of the model parameters under the stable Paretian hypothesis since the index of stability of the innovations is greater than 1 (see Rachev and Mittnik (2000), Tokat and Rachev (2002) and references therein).

After estimating the ARMA(1,1)-GARCH(1,1) parameters, we computed the model residuals and verified which distributional assumption is more appropriate. Figure 1.4 shows a scatter-plot of the estimated (α, β) pairs. For every return distribution in our sample, the estimated index of stability is greater than 1. Even though the estimated values of α are closer to 2 than in the iid model, they are still significantly different from 2.

Comparing the results reported in Table 1.4 to those reported in Table 1.1, we observe that the Gaussian model is rejected in fewer cases in the ARMA-GARCH model than in the simple iid model; nevertheless, the Gaussian assumption is rejected for more than 82% of the stocks at the 99% confidence level using the truncated data. The stable Paretian assumption is rejected in only about 6% of the stocks at the same confidence level (see Table 5).

A summary of the computed statistics for the residuals of the ARMA-GARCH model is reported in Table 1.6. Generally, the results imply that the stable Paretian assumption is more adequate as a probabilistic model for the innovations compared to the Gaussian assumption. As in the iid model, truncating the data improves the AD-statistic in the Gaussian case but still it

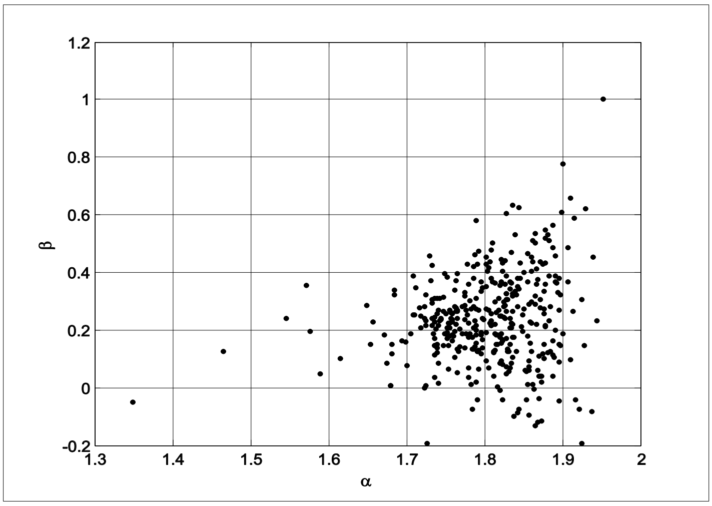

Fig. 1.4. Scatter plot of the index of stability and the skewness parameter for the residuals in the ARMA-GARCH model for all 382 stocks.

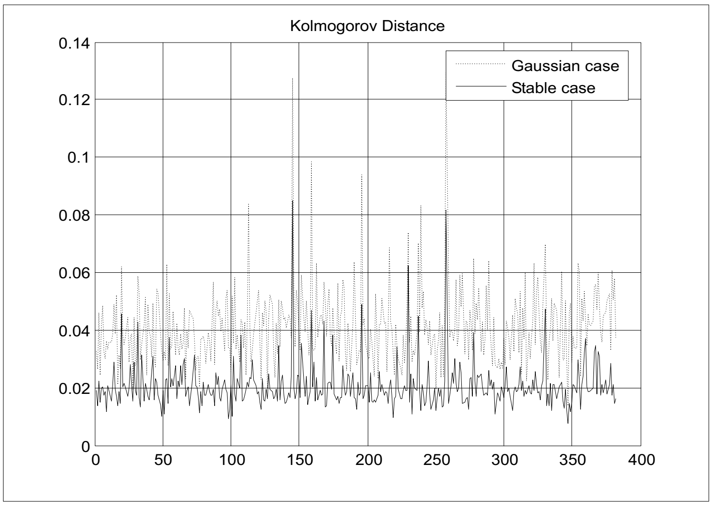

Fig. 1.5. Kolmogorov Distance in both stable Paretian and Gaussian cases for the residuals in the ARMA-GARCH model.

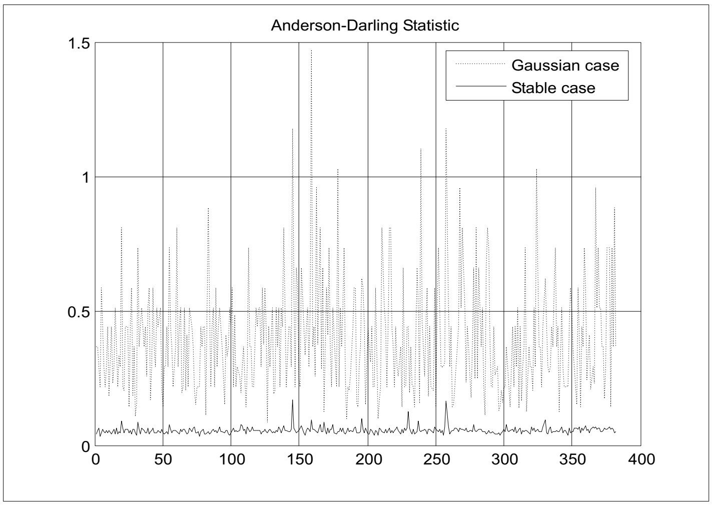

Fig. 1.6. Anderson-Darling statistic in both stable Paretian and Gaussian cases for the residuals in the ARMA-GARCH model.

Confidence level	95%	99%	99.50%	99.90%
Original data	97.38%	92.94%	89.79%	79.58%
Truncated data	93.46%	82.46%	75.92%	61.78%

Table 1.4. Percentage of stocks for which normality is rejected at different confidence levels using Kolmogorov distance in the ARMA-GARCH model.

Confidence level	95%	99%	99.50%	99.90%
Original data	12.05%	5.76%	4.98%	2.62%
Truncated data	13.87%	6.28%	4.50%	3.40%

Table 1.5. Percentage of stocks for which the Stable Paretian hypothesis is rejected at different confidence levels using Kolmogorov distance in the ARMA-GARCH model.

remains about seven times larger than the AD-statistic in the stable Paretian case.

1.5 A back-testing example

We perform a back-testing analysis for the time series of the equity ILN, which is one of the 382 equities considered in the previous analyses. The model parameters are estimated using a moving time window of 1000 observations and the back-testing period is 500 and 1500 days. We compare the performance

		α	β	KD Normal	KD Stable	AD Normal	AD Stable
	mean	1.805	0.24	0.0502	0.0198	0.4859	0.0525
Original	median	1.813	0.23	0.0464	0.0184	0.4389	0.0512
Data	25% quantile	1.763	0.14	0.0381	0.0159	0.365	0.0465
	75% quantile	1.856	0.33	0.0566	0.0219	0.5867	0.0566
	mean	1.834	0.28	0.041	0.0205	0.4061	0.0548
Truncated	median	1.842	0.26	0.0391	0.0192	0.365	0.0533
Data	25% quantile	1.794	0.16	0.033	0.0164	0.2549	0.0474
	75% quantile	1.883	0.39	0.0466	0.0227	0.5128	0.0595

Table 1.6. Summary statistics for the entire sample of 382 stocks, ARMA-GARCH model.

of the more simple GARCH(1,1) model from the ARMA-GARCH family with stable and normal innovations with respect to risk estimation in the cases of Value-at-risk (VaR) and Expected tail loss (ETL) risk measures at 99.5% and 99.9% confidence levels. The performance is compared in terms of the number of exceedances for the VaR measure, that is how many times the forecast of the VaR is above the realized asset return. We verify if the number of exceedances is in the 95% confidence interval for the corresponding back-testing period.

Confidence level	Stable VaR	Normal VaR	95% Confidence interval
99.5% (500 days)	4	10	[0, 5]
99.9% (1500 days)	2	10	[0, 3]

Table 1.7. The number of exceedances of the VaR at a given confidence level in the back-testing period, GARCH(1,1) model.

Confidence level	Stable average ETL	Normal average ETL
99.5% (500 days)	0.1878	0.1231
99.9% (1500 days)	0.3142	0.1433

Table 1.8. he average ETL in the back-testing period of 1500 days produced by both GARCH(1,1) models.

Exactly as expected, the exceedances of the Normal model are not in the 95% confidence interval, in contrast to the exceedances of the Stable model (see Table 1.7). Therefore the Stable model better approximates the tail of the empirical equity return distribution. The average index of stability of the residuals for the entire back-testing period of 1500 days is 1.83.

The ETL produced by the stable model is more conservative than the corresponding figure of the normal model (see Table 1.8). The VaR analysis suggests that the Stable ETL is more realistic than the Normal one.

1.6 Conclusions

We have studied the daily equity returns distribution of 382 U.S. companies comparing the Gaussian and the stable Paretian hypotheses in the context of two assumptions — independent and identical distribution of the daily stock returns and the ARMA-GARCH model. For both models, we strongly reject the normality assumption in favor of the stable Paretian Hypothesis.

References

AKGIRAY, A. and BOOTH, G.G. (1988): The stable-law model of stock returns. *Journal of Business and Economic Statistics, 6, 51–57.*

AKGIRAY, V. and LAMOUREUX, G.G. (1989): Estimation of stable-law parameters: A comparative study. *Journal of Business and Economic Statistics, 7, 85–93.*

BOLLERSLEV, T. (1986): Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31, 307–327.*

DUMOUCHEL, W.H. (1971): Stable Distributions in Statistical Inference. *A dissertation Presented to the Faculty of the Graduate School of Yale University in Candidacy for the Degree of Doctor of Philosophy Yale University.*

ENGLE, R. F, (1982): Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica, 50, 276–287.*

FAMA, E., (1963): Mandelbrot and the stable Paretian hypothesis. *Journal of Business 36, 420–429.*

FAMA, E., (1965): The behavior of stock market prices. *Journal of Business 38, 34–105.*

JANICKI, A. and A. WERON, (1994): *Simulation and Chaotic Behavior of Alpha-Stable Stochastic Processes.* Marcel Dekker, New York.

MANDELBROIT, B., (1963): The variation of certain speculative prices. *Journal of Business 26, 394-419.*

RACHEV, S.T. and S. MITTNIK, S., (2000): *Stable Paretian Models in Finance.* John Wiley & Sons, Chichester.

RACHEV, S.T., (Ed.), (2003): *Handbook of Heavy Tailed Distributions in Finance.* Elsevier/North-Holland, Amsterdam.

SAMORODNITSKY, G. and M. TAQQU, (1994): *Stable Non-Gaussian Random Processes, Stochastic models with Infinite Variance.* Chapman and Hall, New York.