Stable Paretian Distributions: an Unconditional Model for PSI20, DAX and DJIA Indexes

PUBLICADO NA **REVIEW OF FINANCIAL MARKETS**

J. J. Dias Curto
ISCTE: Invited Lecturer
E-mail: dias.curto@iscte.pt

Elizabeth Reis
ISCTE: Associate Professor

José Paulo Esperança
ISCTE: Associate Professor

ABSTRACT

The family of stable distributions is usually regarded as an appropriate probability model for stock returns because of their theoretical properties: only stable distributions have domains of attraction and a stable distribution belongs to its own domain of attraction; and it was empirically observed that stock returns distributions display fatter tails than the normal distribution. In this empirical study we test the returns of Portuguese, German and US stock indexes' conformity to stable Paretian distributions with characteristic exponent bellow 2, by comparison with the normal distribution, which is also stable Paretian but with a characteristic exponent equal to 2.

Keywords: Market efficiency, Portfolio theory, Leptokurtosis, Fat tails, Stable Paretian, Stable-Lévy, Gaussian, Log-likelihood, KD, AIC, AICC and SBC.

ACKNOWLEDGEMENTS

We wish to thank an anonymous referee for constructive comments.

PUBLICADO IN: REVIEW OF FINANCIAL MARKETS, VOL. V-nr 1 2003

I. INTRODUCTION

The assumptions of normality and finite variance of financial asset returns are required by most of the classical financial models, namely the Markowitz's Portfolio Theory, the Capital Asset Pricing Model (CAPM) of Sharpe, Lintner and Mossin and the Black-Scholes' formula.

These models critically depend on the normal form of the probability distribution describing security returns. The Gaussian hypothesis was not seriously questioned until the seminal papers of Mandelbrot (1963) and Fama (1965) were published. Since then, numerous studies have found that the empirical unconditional distribution of returns on financial assets exhibit fatter tails and are more peaked around the center than would be predicted by a Gaussian distribution, i.e., the presence of leptokurtosis in the empirical distributions seems indisputable (Fama, 1965), and the variability of returns seems to be nonstationary.

As a result, alternative distributions possessing such characteristics have been proposed as models for asset returns. Mandelbrot (1963) attempted to capture this excess kurtosis by modeling the distribution of continuously compounded returns as a member of stable-Lévy or stable Paretian distributions, which includes the normal distribution as a special case. Several authors, including Fama (1965), Fama and Roll (1968) and So (1987) have since found support for stable Paretian data generating processes in a wide variety of financial time series.

Other researchers have proposed alternative leptokurtic distributions like the Student's *t* (Blattberg and Gonedes, 1974), the Generalized Error Distribution (Box and Tiao, 1962), the Laplace (or double exponential) and double Weibull (Mittnik and Rachev, 1993), but these distributions do not have the attractive central limit property or divisibility of stable distributions.

Although stable distributions were popular in the 1960's and early 1970's, economists have been reluctant to accept the results of Mandelbrot and others, primarily because of the wealth of statistical techniques available for dealing with normal variables and the relative paucity of such techniques for non-normal stable variables (Fama, 1970). More recently, however, new non-Gaussian approaches to issues like security pricing, portfolio management, risk analysis and empirical analysis had been developed (Rachev and Mittnik, 2000) and the stable Paretian hypothesis enjoyed again quite some popularity among financial modelers.

The existing literature found mixed empirical evidence in favour of the stable Paretian hypothesis. In its seminal paper Fama (1965) concludes that the daily returns of large US mature companies as well as the DJIA stock index follow stable Paretian distributions with characteristic exponent less than 2, which means that the Mandelbrot (1963) hypothesis seems to fit the data better than the Gaussian hypothesis.

In the statistical analysis concerning the daily returns for 30 German stocks forming the DAX share index as well as the DAX itself during the period 1988-1994, Lux (1996) found some evidence in favour of the stable Paretian hypothesis when estimating the parameters of the stable laws and performing standard goodness-of-fit tests. However, the application of a semiparametric technique for analysis of the limiting behaviour in the tails of a distribution (Hill's tail index estimator) suggests that the empirical tail regions are thinner than expected under a stable distribution. Ranging from 2.5% to 15% of all observations, Lux (1996) reports the resulting Hill estimates and concludes that in all cases, $\alpha > 2$, and thus the α –stable hypothesis can be rejected. However, as MucCulloch (1997) and Rachev and Mittnik (2000) argue, the Hill estimator, considering 2.5 and 15 percent of the sample size, will severely overestimate and lead to false conclusions when applied to α –stable data.

In this paper we examine the unconditional distribution of the daily returns on three stock indexes: the Portuguese PSI 20, the German XETRA DAX and the US DJIA, focusing on stable Paretian distributions. The selection of these stock markets was based on their levels of development: the Portuguese market is smaller, more recent and relatively less liquid; the German market is representative of intermediary markets, with a longer history but with a lower role on the economy than typical Anglosaxon capital markets; finally, these markets are compared with the American market, the most liquid and one of the oldest in the World.

Our main purpose is to compare the unconditional fits of both the normal and the stable Paretian distributions, estimating the parameters of the stable laws and performing standard goodness-of-fit tests. The paper is organized as follows. In Section II we present the stable Paretian distributions and their properties. The relationship between the Efficient Market Hypothesis (EMH) and the stable Paretian distributions is presented in section III. Section IV looks at the returns of three stock indexes and considers approaches to assessing the validity of

the normal hypothesis. Section V reports the unconditional fits of the normal and the stable Paretian distributions to the data. The final section summarizes our concluding remarks.

II. THE STABLE PARETIAN DISTRIBUTIONS

Paul Lévy (1924) was probably the first to initiate the investigation of stable distributions. Lévy showed that the tail probabilities approximate those of the Pareto distribution, hence the term “stable Pareto-Lévy” or “stable Paretian” distribution (Campbell, et *al.* 1997). If the right tail of a distribution is asymptotically Pareto then, for large *x*,

$$(1)\quad 1-F(x)\approx cx^{-\alpha_P},$$

where $F(x)$ is the cumulative distribution function, α_P is the tail index and $\alpha_P > 0$, $c > 0$.

The stable Paretian distributions are a natural generalization of the Gaussian distribution. However, non-normal stable distributions have more probability mass in the tail areas than the normal. In fact, the non-normal stable distributions are so fat-tailed that their variance and all higher moments are infinite (Campbell, et *al.* 1997).

Beyond the normal distribution, the Cauchy distribution, the Lévy distribution, and the reflection of the Lévy distribution, there are no closed form expressions for generally stable densities. Instead, the stable Paretian distributions can be expressed by their characteristic function and the most common parameterization is:

$$(2)\quad E\left(e^{itX}\right)=\begin{cases}\exp\left\{i\mu t-\sigma|t|^{\alpha}\left[1-i\beta\tan\dfrac{\pi\alpha}{2}\operatorname{sign}(t)\right]\right\} & \text{if}\quad \alpha\neq 1\\ \exp\left\{i\mu t-\sigma|t|\left[1+i\beta(2/\pi)\operatorname{sign}(t)\ln|t|\right]\right\} & \text{if}\quad \alpha=1\end{cases},$$

where $\operatorname{sign}(t)=\begin{cases}-1 & \text{if } t<0\\ 0 & \text{if } t=0\\ 1 & \text{if } t>0\end{cases}$.

Stable Paretian distributions have four parameters: a location parameter (μ), a scale parameter (σ), an index of skewness (β) and a measure of the height of the extreme tail areas of the distribution, the characteristic exponent (or the stability index) α.

The characteristic exponent of a stable Paretian distribution α determines the total probability in the extreme tails of the distribution and can take any value in the interval $0 < \alpha \leq 2$. When $\alpha = 2$, the relevant stable Paretian distribution is the normal distribution with mean μ and variance $2\sigma^2$. As α decreases from 2 to 0, the tail areas of the stable distribution became increasingly fatter than the normal. For $\alpha < 2$, the sth absolute moment exists only for $s < \alpha$. Thus, except for the normal case ($\alpha = 2$), the stable Paretian distributions has infinite variance.

The parameter β can take any value in the interval $-1 \leq \beta \leq 1$. When $\beta = 0$, the distribution is symmetric around μ. If $\alpha \neq 1$, for positive (negative) β, the distribution is skewed to the right (left). The direction of skewness is reversed if $\alpha = 1$.

μ is the location parameter and can take any value in the interval $-\infty < \mu < +\infty$. When $1 < \alpha < 2$ the stable Paretian distribution has a finite mean given by μ. For $0 < \alpha \leq 1$ the tails are so heavy that even the mean does not exist. In this case, μ should be another parameter (the median, for example, when $\beta = 0$).

The scale parameter σ can take any positive value: $0 < \sigma < +\infty$. When $\alpha < 2$ the variance of a stable Paretian distribution does not exist and σ defines the scale of the distribution, but it will be not the variance. For example, when $\alpha = 1$ and $\beta = 0$ (Cauchy distribution) σ is the semi-interquartile range.

There are several reasons for using a stable distribution to describe financial data. First, only stable distributions have domains of attraction: the Generalized Central Limit Theorem of Feller (1966), hereafter GCLT, establishes that if the distribution of a sum of independent and identically distributed random variables exists, then it must be a member of the stable Paretian class of distributions. When the variance is finite, by the Central Limit Theorem, hereafter CLT, the limiting distribution is the normal distribution; otherwise, when the variance is infinite, the limiting distribution is the stable Paretian non-Gaussian. Second, the stability and invariance under addition: the distribution of sum of independent, identically distributed, stable Paretian variables is itself stable Paretian and, except for origin and scale, has the same form as the distribution of individual summands. It means that a stable distribution belongs to its own domain of attraction.

The third argument for modelling with stable distributions is empirical: many large data sets exhibit heavy tails and skewness.

III. THE STABLE PARETIAN DISTRIBUTIONS AND THE EFFICIENT MARKET HYPOTHESIS

The Efficient Market Hypothesis (EMH) states that prices of securities at every point in time represent the best estimates of its intrinsic values. This implies in turn that, when an intrinsic value changes, the actual price will adjust instantaneously.

The random walk version of the EMH (Bachelier (1900), Osborne (1959) and Fama (1965)), assumes that successive price changes are independent, identically distributed random variables. If the number of transactions per day, week or month is very large, then price changes across these differencing intervals will be sums of many independent variables and by the GCLT might be well approximated by a member of the stable Paretian class of distributions. If the second moments are finite, the CLT leads us to expect that the daily, weekly and monthly returns each have normal or Gaussian distributions otherwise the limiting distribution of low frequency returns is the stable Paretian non-Gaussian.

Therefore, the assumption that returns are normally distributed is not necessarily implied by market efficiency even in the most restrictive random walk version.

If the true generating process of stock returns is stable Paretian with $\alpha < 2$, the price of a security will often tend to jump up or down by very large amounts during very short time periods, i.e., the path of the price level will usually be discontinuous. As a result, the risk is higher because the variability and the probability of large losses are greater when compared with the Gaussian generating process.

If the EMH is true, the stable Paretian distribution with $\alpha < 2$ would imply that intrinsic values often change by large amounts during shorts periods of time in order to justify large and often price changes — a situation that is consistent with a dynamic economy in a world of uncertainty (Fama, 1965). Therefore, both hypothesis are not incompatible.

IV. DATA AND DESCRIPTIVES STATISTICS

The data consists of daily closing prices of PSI20, DAX, and DJIA indexes, which are market indexes for the Portuguese, the German and the US equity markets. These series runs from 31 December 1992 to 31 December 2001, yielding 2227, 2291 and 2269 observations, respectively. The rates of return (not adjusted for dividends) are calculated by taking the first differences of the logarithm of the series:

(3) $r_t = 100\text{x}\left[\ln\left(P_t\right) - \ln\left(P_{t-1}\right)\right]$, where P_t is the closing value for each index at time *t*.

The levels and percent returns are shown in the next figure.

Figure 1 somewhere here.

Table 1 summarizes the basic statistical properties of the data. The mean returns are all positive but close to zero. The returns appear to be somewhat asymmetric, as reflected by negative skewness estimates: more observations are in the left-hand (negative) tail than in the right-hand tail. All three series returns have very heavy tails and show strong departure from normality (skewness and kurtosis coefficients are all statistically different from those of the standard normal distribution which are 0 and 3, respectively). The Jarque-Bera test also clearly rejects the null hypothesis of normality.

Table 1 somewhere here.

According the Ljung-Box statistic for returns, there is no relevant autocorrelation for the DAX and the DJIA indexes, which seems to confirm that the logarithm of prices follows a martingale[1].

For PSI20, returns have statistically significant first order autocorrelation. The successive returns correlation must be due to the market thinness and nonsynchronous trading that is common in many small capital markets and seems to violate the martingale and the efficient market hypothesis in the Portuguese stock market.

However, as stated by Fama (1970), with samples of the size underlying table 1 statistically "significant" deviations from zero covariance is not necessarily a basis for rejecting the efficient

market hypothesis, because it is unlikely that the small absolute levels of serial correlation (0,189) can be used as the basis of substantially profitable trading systems.

V. MODELLING THE UNCONDITIONAL DISTRIBUTION OF RETURNS

To compare the goodness-of-fit of the candidate distributions, the normal and the stable Paretian, we will use the following methodology. First the four parameters of stable Paretian distributions are estimated by the Maximum Likelihood (ML) estimation method using directly the release 3.04 of STABLE software[2].

Then we will adopt the methodology of Mittnik and Paolella (2002) considering four goodness-of-fit criteria. The first is the maximum log-likelihood value obtained from the ML estimation. It may be viewed as an overall measure of goodness-of-fit and allows us to judge which candidate is more likely to have generated the data. The distribution with an upper value of log-likelihood is judged to be preferable.

The second criteria is the Kolmogorov distance:

$$(4)\ KD = 100 \times \sup_{x \in R} \left| F_s(x) - \hat{F}(x) \right|,$$

where $\hat{F}(x)$ denotes the cumulative distribution function of the estimated parametric density and $F_s(x)$ is the empirical sample distribution, $F_s(x) = n^{-1} \sum_{t=1}^{n} I_{(-\infty;x]}\left(\frac{r_t - \hat{\mu}}{\hat{\sigma}_t} \right)$, where $I_{(-\infty;x]}$ is the usual indicator function. The distribution with a lower value of the Kolmogorov distance is judged to be preferable.

The other two are likelihood-based criteria that penalize the complexity of the models and the penalty factor depends on the number of parameters that have been estimated. The first is the bias-corrected Akaike (Hurvich and Tsai (1989); see also Brockwell and Davis (1991)) information criterion (1973) given by:

$$(5)\ AICC = -2 \log L\left(\hat{\theta}\right) + \frac{2n(k+1)}{n-k-2}.$$

The second is the Schwarz (1978) Bayesian criterion defined as,

(6) $SBC = -2\log L\left(\hat{\theta}\right) + \frac{k\ln(n)}{n}$,

where $\log L\left(\hat{\theta}\right)$ is the maximum log-likelihood value, n is the number of observations and k is the number of parameters. The distribution with a lower value for these information criteria is judged to be preferable.

The ML estimates of the fitted unconditional distributions are shown in table 2.

Table 2 somewhere here.

The estimates of the characteristic exponent are statistically bellow 2, the value of the normal distribution, suggesting that the normal hypothesis is innapropriate for all three stock indexes returns when compared to the non-normal stable Paretian distributions. This is supported by the goodness-of-fit measures reported in Table 3.

Table 3 somewhere here.

The stable Paretian distribution is preferred quite strongly with respect to the various goodness-of-fit criteria that have been considered for all three stock market indexes: the maximum log-likelihood value is higher and the KS distance and the information criteria are lower for the stable Paretian distribution. The LR test (in the last row of the table) also clearly rejects the null hypothesis of normality in favour of the non-Gaussian stable Paretian distributions.

These results are consistent with those of Fama (1965) and Lux (1996) for DJIA and DAX stock indexes, respectively, and the normal distribution seems not to be an adequate model to describe the daily returns of the three stock indexes: a stable Paretian distribution with $\alpha < 2$ seems to fit the data better than the Gaussian distribution. Therefore, new non-Gaussian approaches are needed to issues like security pricing, portfolio management, risk analysis and empirical analysis (Rachev and Mittnik, 2000).

We conclude the analysis of table 2 comparing the estimates of stable laws. First, as we referred before, α estimates are all clearly below 2, the value of the Gaussian distribution. Second, the estimates of β are not significant for PSI20 and DJIA (at the 5% level). It means that, if the stable Paretian model is valid, the null hypothesis of symmetry can not be rejected for

both the Portuguese and American stock indexes returns. On the other hand, the DAX returns distribution is statistically skewd to the left. Third, scale estimates are also very similar for the Portuguese and American stock indexes and the volatility is higher in the German DAX. Finally, the expected return is also higher in the German stock market, according to the location estimate.

VI. CONCLUSIONS

Until the seminal paper of Mandelbrot (1963) it was generally accepted that the empirical distributions of stock returns were approximately normal or Gaussian. Since then, a consensus exists that distributions of stocks returns display characteristics that contradict the normality hypothesis.

The rates of return of the Portuguese PSI20, German DAX and US DJIA stock indexes were analyzed in this study, and it has been observed too much mass near the mean and in the extreme tails of the empirical distributions, i.e., there is an excess of kurtosis when compared to the normal distribution, for all three stock indexes returns.

To model this empirical characteristic of returns, it was considered the alternative stable Paretian distributions and the main conclusion of this study is, according to the goodness-of-fit statistics, that the empirical distributions of stock indexes returns conform better to stable Paretian distributions than to the normal distribution. It does not mean, however, that the true generating process of returns must be stable Paretian non-Gaussian!

If the stock returns generating process is stable Paretian with $\alpha < 2$, the EMH is more likely to occur in a dynamic economy, i.e., the large changes in prices must be due to often large changes in the intrinsic values of securities. Such situation is only consistent with a dynamic economy where the securities' analysts can permanently alter price targets according to the news about companies. The Portuguese stock market lower efficiency, as compared to both the German and the US stock markets, is probably due to lower activity by fundamental analysts and slower diffusion of relevant information on stock prices.

REFERENCES

Akaike, H. (1978), *Time series analysis and control through parametric models*, Applied Time Series Analysis, D.F. Findley (ed.), Academic Press, New York;

Bachelier, Louis (1900 and 1964), *Theory of Speculation*, in P. Cootner, ed., The random character of stock market prices, Cambridge, MA: MIT Press;

Black, Fischer e **Scholes**, Myron (1973), *The pricing of options and corporate liabilities*, Journal of Political Finance, May-June.

Blattberg, R. C. e **Gonedes**, N. J. (1974), A comparison of the Stable and Student distributions as statistical methods for stock prices, Journal of Business, 47, pp. 244-280;

Box, G. E. P. e **Tiao**, G. C. (1973), *Bayesian inference in statistical analysis*, Addison-Wesley, Reading, Massachusetts;

Brockwell, Peter J. e **Davis**, Richard A. (1991), *Time series: theory and methods*, Springer-Verlag, 2nd edition, New York;

Campbell, John Y., **Lo**, Andrew W., **MacKinlay**. A. Craig (1997), *The econometrics of financial markets*, Princeton University Press, New Jersey;

Fama, Eugene F. (1965), *The behavior of stock prices*, Journal of Business, XXXVIII (1), January, pp. 34-105;

Fama, Eugene F. e **Roll**, Richard (1968), *Some properties of symmetric stable distributions*, Journal of The American Statistical Association, vol. 63, No 323, September, pp. 817-37;

Fama, Eugene F. (1991), *Efficient capital markets: II*, Journal of Finance, XLVI, 5, December, 1575-1617;

Feller, W. (1966), *An introduction to probability theory and its applications*, John Wiley & Sons, New York;

Hurvich, C. M., **Tsai**, C. L. (1989), *Regression and time series model selection in small samples*, Biometrika, 76, 297-307;

Jarque, C. M. e **Bera**, A. K. (1987), *A test for normality of observations and regression residuals*, International Statistics Review, 55 (2), pp 163-172;

Lévy, Paul (1924), *Théorie des erreurs. La Loi de Gauss et les Lois Exceptionelles*, Bull. Soc. Math., 52 pp 49-85;

Lintner, John (1965), *Security prices, risk and maximal gains from diversification*, Journal of Finance, 20, pp. 587-615;

Mandelbrodt, Benoit B. (1963), *The variation of certain speculative prices*, The Journal of Business, 36, pp. 394-419;

Markowitz, H. M. (1952), *Portfolio Selection*, Journal of Finance (December).

McCulloch, John H. (1997), *Measuring tail thickness to estimate the stable index: a critique*, Journal of Business and Economic Statistics, Vol. 15, No. 1, pp. 74-81;

Mittnik, Stefan e **Rachev**, Svetlozar T. (1993), *Modeling asset returns with alternative stable models*, Econometric Reviews 12, pp. 261-330;

Mittnik, Stefan, **Paolella**, Marc S. e **Rachev**, Svetlozar T. (1998), *Unconditional and conditional distributional models for the Nikkei Index*, Asia-Pacific Financial Markets 5, pp. 99-1281;

Mittnik, Stefan e **Paolella**, Marc S. (2002), *Prediction of financial downside-risk with heavy-tailed conditional distributions*, in Handbook of heavy tailed distributions in finance, North-Holland;

Mossin, Jan (1968), *Optimal multiperiod portfolio policies*, Journal of Business, 41, pp.215-229;

Osborne, M. F. M. (1959), *Brownian motion in the stock market*, Operations Research, 7, pp. 145-173;

Paolella, Marc S. (2001), *Testing the stable Paretian assumption*, Mathematical and computer modeling, 34, pp 1095-1112;

Rachev, Svetlozar e **Mittnik**, Stefan (2000), *Stable Paretian models in finance*, John Wiley & Sons, New York;

Schwarz, G. (1978), *Estimating the dimension of a model*, Annals of Statistics, 6, pp. 461-64;

Sharpe, William. F. (1964), *A theory of market equilibrium under conditions of risk*, Journal of Finance, 19, pp. 425-442;

So, J. C. (1987), *The sub-Gaussian distribution of currency futures: stable Paretian or non stationarity*, Review of Economics and Statistics, pp. 100-107.

Figure 1: Levels and percent returns

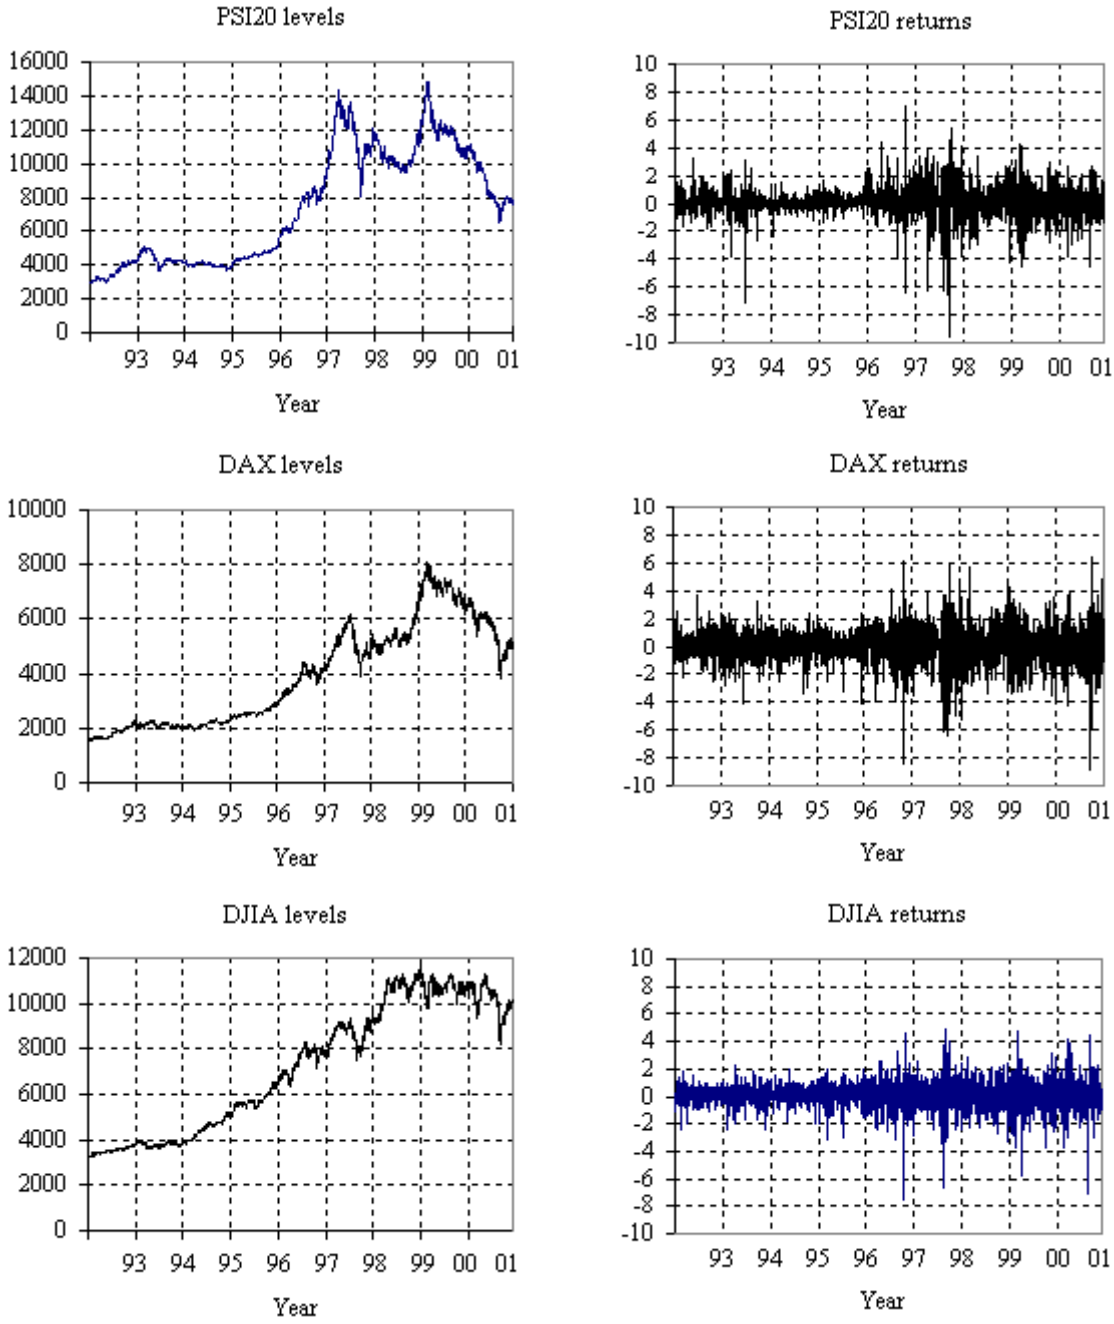

Table 1: Summary statistics of returns

Statistics	PSI20	DAX	DJIA
No. of observations	2226	2290	2268
Mean	0.043106	0.053001	0.048963
Median	0.038840	0.106379	0.066329
Maximum	6.941259	6.415896	4.860544
Minimum	-9.589772	-8.823020	-7.454935
Standard deviation	1.217006	1.830579	1.011597
Skewness	-0.663949	-0.482210	-0.532190
Kurtosis	10.621	6.422	8.242
JB[a]	5551.07*	1206.32*	2703.53*
LB(10)[b]	94.87*	15.76	17.33
$\hat{\rho}_1$ [c]	0.189*	0.013	0.023
$\hat{\rho}_2$ [c]	0.005	-0.041	-0.04
$\hat{\rho}_3$ [c]	0.033	0.005	-0.018

*Denotes significant at the 1% level;
[a]JB is the Jarque–Bera test for normality;
[b]LB(10) is the Ljung Box test for returns;
[c] $\hat{\rho}_i$ are the estimates of autocorrelation coefficients for returns;

Table 2: Estimates of the stable Paretian distribution[a]

Stock Indexes	α	β	σ	μ
PSI20	1.5586	0.0512	0.5504	0.0446
	(0.0336)	(0.0703)	(0.0122)	(0.0205)
DAX	1.7559	-0.2958	0.8080	0.1348
	(0.0299)	(0.1010)	(0.0156)	(0.0296)
DJIA	1.7133	-0.1716	0.5763	0.0922
	(0.0314)	(0.0918)	(0.0115)	(0.0213)

[a] Standard errors are in parentheses.

Table 3: Goodness-of-fit of unconditional distributions

	PSI20		DAX		DJIA	
Statistics	Normal	Stable	Normal	Stable	Normal	Stable
L[a]	-3376.64	-3128.86	-3941.17	-3836.70	-3230.73	-3086.83
KD[b]	8.38	1.90	4.10	1.69	5.44	1.78
AICC[c]	6755.28	6261.72	7884.34	7677.40	6463.57	6177.66
SBC[d]	6753.29	6257.73	7882.35	7673.41	6461.47	6173.67
LR[e]	495.56*		1211.94*		6173.66*	

*Denotes significant at the 1% level;
[a]L refers to the maximum log-likelihood value, [b]KD is the Kolmogorov Distance, [c]AICC is the bias-corrected Akaike Information Criterion, [d]SBC is the Schwarz Bayesian Criterion and [e]LR is the likelihood ratio test for two restrictions: $\alpha = 2$ and $\beta = 0$.

[1] The practice in the efficient capital markets literature is to speak of stock prices as following a martingale. In such cases "price" should be understood to include reinvested dividends (LeRoy, 1989).

[2] It is available on the *site* of professor John Nolan: http://www.cas.american.edu/~jpnolan.