Stable GARCH and Temporal Aggregation*

Zoltán Palágyi
Budapest University of Economics, and Institute of Economics.
E-mail: palagyi@bkae.hu

Gábor Kőrösi
Institute of Economics, Budapest, and CEU.
E-mail: korosi@econ.core.hu

November 2002

Abstract

The paper presents estimates of stable power GARCH models for returns aggregated with varying temporal frequency. Although the stable GARCH model gradually converges towards a GARCH model based on normally distributed innovations, convergence is slow, and stable GARCH clearly dominates 'standard' (power) GARCH.

When innovations follow a Lévy stable distribution, the properties of standard tests and estimators may change substantially. We explore the behaviour of the standard portmanteau test under stable distribution.

The empirical work is based on two substantially different investment. We found substantial difference between the properties of a very liquid paper traded on a sophisticated market, and one which is bought by much fewer investors on a thin market. Stable GARCH model seems to be more relevant in emerging capital markets, indicating that investors are more likely to be hit by extreme shocks on an emerging capital market.

Keywords: Lévy distribution, Stable GARCH, Aggregation, Risk modelling, Emerging capital markets.
JEL classification: C22, G12.

*This research was supported by a grant from the CERGE-EI Foundation under a program of the Global Development Network. All opinions expressed are those of the author(s) and have not been endorsed by CERGE-EI, or the GDN. We received valuable comments from Temisan Agbeyegbe and Jan Hanousek to an earlier version of this paper.

1 Introduction

GARCH models are the single most important econometric tools for describing the return process of securities.[1] These models reasonably depict some typical characteristics of the return series recorded at exchange markets (e.g. volatility clustering, heavy tails), but they still fail in many cases. The proliferation of alternative GARCH models[2] indicates the intensity of the effort to improve the performance of these models.

A possible reason for the failures of GARCH may be the assumption about the distribution of innovations: diverse GARCH variants frequently assume normally distributed innovations,[3] however, that is often not supported by empirical evidence. On the one hand, we can find examples of return series, for which GARCH filtered residuals are skewed and/or leptokurtic—a behavior inconsistent with the normal assumption. Many recent applications therefore use more heavy tailed distributions for the innovation process, e.g., Student-t with low degrees of freedom. However, the usage of these distributions seems to be arbitrary, lacking theoretical justification. On the other hand, an interesting example is given by Aït Sahalia et al. (2001), who find substantial discrepancies between the option implied and the asset implied state-price densities for the S&P 500 index of the NYSE, which, under some regularity conditions and the no-arbitrage assumption should be identical. They attribute the differences to a Poisson jump process generating excess skewness and kurtosis in the unconditional distribution of asset prices. Jumps correspond to recurring extreme events, the effect of which may not be captured by a GARCH process driven by innovations from a distribution with finite kurtosis (e.g., Gaussian or Student-$t(k)$ with $k > 4$). Modelling these extreme events is, however, very important for practical problems like assessment of the risk of an investment. Non-parametric methods are unable to deal with relatively rare events; thus we need a parametric approach.

Choosing a distribution to model GARCH innovations is still an open problem. None of the candidates tried so far (e.g. normal, Student-t, generalized

[1] GARCH is widely used among practitioners: for example, the most popular VaR (Value at Risk) model is based on an integrated GARCH model, with predetermined parameters.

[2] Bollerslev et al. (1994) give a long list of alternative models of the family.

[3] In practice, quasi ML is usually used, in which case the underlying distribution can be different from normal, but it should still maintain the most important characteristics of a Gaussian innovation process.

exponential, Pareto stable distributions) have proven to be significantly better in general than the others. Stable distributions (normal is a special case) can be favoured against the others by their special role in probability theory: by the generalized central limit theorem. Stable distributions are the only non-degenerate distributions arising as limits of normalized sums of independent and identically distributed random variables. Thus if we think of innovations as sums of random effects too numerous and difficult to incorporate into the model, then stable distributions are a natural choice to describe them. The name stable refers to stability under addition: the distribution of appropriately normalized sums of iid stable distributions is the same as the distribution of the summands. The key parameter of stable distributions is the index of stability (this parameter is invariant under convolution) $0 < \alpha \leq 2$. For the normal distribution $\alpha = 2$. Besides α there is a location, a scale and a skewness parameter, thus stable distributions are rather flexible.[4]

Both the so called Pareto stable models (returns are considered iid, following a stable law with $\alpha < 2$) and GARCH models with normal or Student-t innovations are able to describe the heavy tail property of return series, and in this sense these models have been considered as competing models by several authors. The question that motivated some studies was whether the fat tail property of data is "produced" by infinite variance stable distributions, or GARCH driven by finite variance innovations. Several authors are reluctant to accept infinite variance random variables as building blocks for econometric models. The results of some studies make it clear though, that the above models cannot be compared. Ghose and Kroner 1995 show for example that data simulated from some properly parameterized integrated GARCH models seem to have the stability under summation property, and one might erroneously conclude that they actually come from a stable distribution. Thus, it only seems to make sense to check for the stability of GARCH filtered innovations.

The use of stable innovations in GARCH models has been proposed by several authors: McCulloch (1985), Liu and Brorsen (1995), Panorska, Mittnik and Rachev (1995), Mittnik, Paolella and Rachev (2000), but they are still rarely used by practitioners, mainly because of theoretical problems (identification, stationarity conditions, etc.), and computational difficulties. Some problems have been

[4]The $t(1)$ distribution (the Cauchy distribution) is the symmetrical stable distribution with $\alpha = 1$.

solved, and some are topics of active research.

We have a dual goal in this paper. On the one hand, we compare the performance of GARCH models with Gaussian and Lévy distributions (stable with $\alpha < 2$). We use return series with different frequencies, to analyse the stability of the risk process. On the other hand, we compare the properties of two investments at two very different markets: one major stock from NASDAQ, and one from the Budapest Stock Exchange (BSE), which is a recently (re)established, small, thin market with little tradition and experience in the proper management of an exchange market.

Section 2 describes some important market characteristics, and preliminary analysis of the data. Section 3 outlines major features of stable GARCH processes, and empirical results, while Section 4 concludes.

2 Markets and Data

Our samples come from transaction level data from 1998, full year. We consider returns of two stocks, CISCO traded at NASDAQ, and MOL traded at BSE.

There is probably no need to present the operation of NASDAQ, or to describe one of its most heavily traded stocks of the period, CISCO. We chose this stock because it is well known, many investors traded in it, and it had such a high liquidity, that the properties of its return process may be expected to reflect the characteristics of an efficient capital market.

The Budapest Stock Exchange was gradually set up during the period 1988-90 as a bond market, and it was inaugurated in 1990. Regular daily trading started early 1991. Initially, there was only one share traded on the floor, and the number of listed companies increased very slowly. As the market was extremely thin, little trading occurred after the novelty faded. It was a very small, negligible, lacklustre market until after the 1995 macroeconomic stabilization package. The market, however, was set up by professionals keen on creating a properly functioning modern stock exchange. Thus, it probably was the best-regulated and most transparent market in the Central and Eastern European region.

Hungary chose a more gradualist approach to economic transformation than most other transition economies. The careful constitution of the institutional framework of a market economy got strong preference. By the mid-1990's it could develop into a 'normal' market, leaving behind most of the initial peculiarities. Both market participants and regulatory authorities acquired the skills necessary for operating smoothly on the market (c.f., Johnson and Schleifer, 1999).

The gradualism of the Hungarian transition was not restricted to institution building. Despite some timid attempts for mass privatisation, most former state-owned enterprises were sold on a case-by-case basis to (usually strategic, frequently foreign) investors, who typically got majority stakes. Privatisation through the stock exchange was rarely used, although later several large privatised companies were introduced to BSE, usually well after the majority was sold to a strategic investor. Thus, some of the biggest Hungarian enterprises are in fact traded on BSE.

MOL (Hungarian Oil Co) is one of the largest privatised companies listed on BSE. Unlike most other large Hungarian corporations, MOL has no single dominant owner, thus its share price is clearly determined by market forces. However, MOL is large on the Hungarian market only; it is a very small company compared to the multinational firms in the oil sector, or, indeed, compared to CISCO. Although many Hungarians invested into the stock exchange, and into MOL after the boom started in late 1996, the bulk of trading was (and still is) executed by a handful of foreign investment funds, thus it is traded on an oligopolistic market. So we expected the return series of the two stocks to have rather different characteristics. (c.f., Palágyi and Mantegna (1999) and Palágyi et al. (2001).)

Our sources for transaction level data were TAQ (CISCO), and BSE information office (MOL). Trading time was recorded with the precision of one second in both markets. Transactions registered at the same second were consolidated as the first step of data processing. A well-known characteristic of transaction-level data is that prices change relatively infrequently (Campbell et al., 1997). In our case, more than half of transaction level returns were zero for any largish subsample; the ratio for MOL was as large as 2/3.

For a preliminary analysis we described transaction-level returns as a mixture of a Lévy and a degenerate (constant zero) distribution. In this approach some of the zero returns may be considered as coming from the Lévy distribution, but since we cannot identify these zeros, we simply left them out of the estimation of the index of stability (α). This way (as our Monte Carlo simulations showed) we slightly overestimated the value of α, however, the upward bias was of the order of the standard error of estimation.

The question whether a series of returns $r_1, ..., r_m$ is stable under addition can be investigated by forming non-overlapping sums of size n of successive returns $r_1^{(n)} = r_1 + ... + r_n$, $r_2^{(n)} = r_{n+1} + ... + r_{2n}, \ldots$, and estimating the index of stability α_n of the sequence $r_j^{(n)}$ of aggregated returns. If the original sequence of returns is stable, then α_n is a constant independent of n. If we start with non-zero transaction level returns, then $r_j^{(n)}$ can be interpreted as the return of the asset after n successive price changes on the market. On the other hand, if we start with one-minute returns, $r_j^{(n)}$ is the return on the asset after n minutes.

Figure 1. depicts the convergence patterns of α_n for CISCO and MOL, using both transaction-level, and one-minute returns as initial (disaggregated) series. To facilitate a proper comparison of the graphs we note that on the average 964/160 non-zero price changes occur in one hour for CISCO/MOL respectively. While alpha estimated from 60 minute returns of CISCO is about 1.7 (top right panel), alpha estimated from the returns of CISCO after 964 successive price changes is 2 (top left panel). On the left-hand side graph alpha converges rapidly to two, while on the right one it converges very slowly. In fact, stability of one minute returns of CISCO might not be rejected by a formal test. The explanation of this phenomenon most probably is that returns over equidistant time intervals are defined by (a random number of) non-zero transaction level returns falling into the interval. Thus, we may think of the distribution of say 30 minute returns of CISCO as a mixture of distributions with different parameters.

Pictures for MOL (bottom panels) look very different at first glance, but we must take into account the fact that the market for CISCO is about 6 times faster (in terms of number of price changes) than that for MOL. In fact the left graph of MOL looks a bit like a magnified version of the left graph of CISCO. Using the *1 hour=160 price changes* conversion the right graph of MOL can be considered the same as the left pushed down by 0.2. The right graph gets very noisy at the end, because sample size gets smaller with aggregation, however, fluctuations can be interpreted as the size of error, and then the end of the right graph is on the average a horizontal line.

Based on the above results it seems reasonable to try to model returns over fixed physical time horizons. We use stable GARCH in this paper.

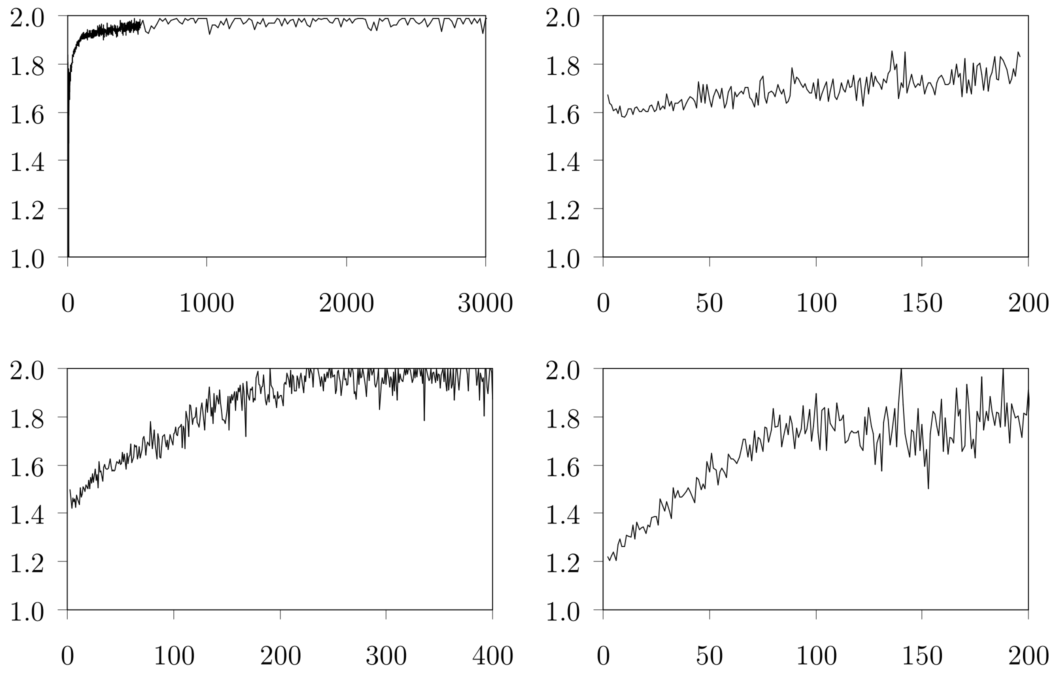

Figure 1: Convergence patterns of the index of stability (α) of returns for the stocks CISCO (top panels), and MOL (bottom panels) as a function of the level of aggregation (n). On the left panels non-zero transaction level returns are aggregated while on the right panels one-minute returns are aggregated. The average number of price changes per hour is 964 for CISCO, and 160 for MOL.

3 Stable GARCH estimation results

The series y_t is said to follow a stable power GARCH, $S^{\delta}_{\alpha,\beta}\ GARCH(r,s)$ process (Mittnik, Paolella, Rachev (2000)) if

$$y_t = \mu_t + \epsilon_t,$$

where $\epsilon_t = \sigma_t z_t$, $z_t \sim S_\alpha(1,\beta,0)$, and

$$\sigma_t^{\delta} = c_0 + \sum_{i=1}^{r} c_i |\epsilon_{t-i}|^{\delta} + \sum_{j=1}^{s} d_j \sigma_{t-j}^{\delta}\,.$$

$S_\alpha(\sigma,\beta,\mu)$ denotes the stable distribution with parameters α (index of stability), σ (scale), β (skewness), μ (location). We use the parametrization of Samorodnitsky and Taqqu (1994). In our analysis we use symmetric distributions only

($\beta = 0$), μ_t denotes a strictly stationary ARMA process.

Sufficient conditions for the strict stationarity of the $S^{\delta}_{\alpha,\beta}$ $GARCH(r,s)$ process with $1 < \alpha \leq 2$ and $0 < \delta < \alpha$ are given in Mittnik, Paolella, Rachev 2002: $c_0 > 0$, $c_i \geq 0$, $i = 1, ..., r$, $r \geq 1$, $d_j \geq 0$, $j = 1, ..., s$, $s \geq 0$, $r \geq s$, and

$$V = E|z_t|^{\delta} \sum_{i=1}^{r} c_i + \sum_{j=1}^{s} d_j \leq 1.$$

If $1 < \alpha \leq 2$ and $0 < \delta < \alpha$, $E|z_t|^{\delta}$ can be written in the following closed form:

$$E|z_t|^{\delta} = \lambda_{\alpha,\beta,\delta} = \frac{1}{\psi_{\delta}} \Gamma(1 - \frac{\delta}{\alpha})(1 + \tau^2_{\alpha,\beta})^{\frac{\delta}{2\alpha}} \cos(\frac{\delta}{\alpha} \arctan \tau_{\alpha,\beta}),$$

where $\tau_{\alpha,\beta} := \beta \tan(\alpha\pi/2)$ and

$$\psi_{\delta} = \begin{cases} \Gamma(1-\delta) \cos \frac{\pi\delta}{2}, & \text{if } \delta \neq 1, \\ \pi/2, & \text{if } \delta = 1. \end{cases}$$

If $\alpha < 2$ and $\delta \geq \alpha$, then $E|z_t|^{\delta} = \infty$ (if $\alpha < 2$ and $\delta \to \alpha$, then $\lambda_{\alpha,\beta,\delta} \to \infty$). Mittnik, Paolella and Rachev (2002) examine the case $\delta = \alpha < 2$ by Monte Carlo simulations, and conclude that in this case the process is not stationary.

The above model incorporates the 'usual' GARCH (Bollerslev, 1986) as a special case when $\alpha = \delta = 2$ (one can easily check that at the limit $\alpha \to 2$, $\delta \to 2$ the above stationarity condition on V becomes the 'usual' stationarity condition for GARCH), and the symmetric models in the APARCH family of Ding et al. (1993), when $0 < \delta < \alpha = 2$.

We estimated the above model by Maximum Likelihood, using the Nelder-Mead polytope method of Press et al (1992). Maximization of the log likelihood function was restricted to the parameter space satisfying the stationarity conditions. As there is no closed form of the relevant partial derivatives of the likelihood function, standard errors and confidence intervals had to be estimated by the bootstrap method. (LePage et al., 2001)

Before estimating the model, we tested for autocorrelation in the series of returns and squared returns. We calculated the most popular portmanteau statistics (Ljung-Box) from the returns series. Unfortunately, the distribution of the test statistics is unknown under the null hypothesis that the data generating process follows an i.i.d. Lévy stable distribution. If the moments of the data generating process are finite, at least up to the fourth order, its asymptotic distribution is χ^2. However, non-Gaussian stable distributions do not have finite second and higher moments. Thus, the asymptotic distribution of the Ljung-Box statistics is unknown, and it may not exist. Therefore, we tabulated the response surface of the Ljung-Box statistics for various values of α in Table 1, assuming 2000 observations, using 100000 repetitions.[5] Further, as the moments of the data generating process do not exist, there is no guarantee that the Ljung-Box statistic has the same critical values if it is computed from the powers of the series. Table 2 gives critical values when the statistic is computed from the autocorrelations of the 'δ^{th}' power of the observations, where $\delta = \alpha - 0.001$,[6] while Table 3 tabulates the critical values for the squared time series.

Indeed, our expectation that the properties of these three test statistics are different is justified. Critical values corresponding to the same α, and order of the autocorrelation test are frequently substantially different. As α decreases the distribution becomes more extreme: while the 95% ordinate of the distribution for the Ljung-Box test statistic seems to converge towards 0 with decreasing α and increasing power, the 99^{th} ordinate tends to increase.

We also did some sensitivity analysis with respect to the sample size. It appears that some very extreme values may emerge as sample-size increases. The distribution seems to gradually degenerate to a small number of extremely large test statistics, while most of the values become very small. Thus the normalisation of the distribution of the portmanteau test may not be correct when innovations follow stable distribution. Which really means that we have to simulate the critical values for each sample size separately.

When testing the autocorrelation of the relevant powers of returns, thus we simulated critical values directly corresponding to the actual sample information (see Table **??**). We first estimated the index of stability (α) from the return series, then generated 1000 samples of iid α stable random numbers. The size of the simulated samples agreed with the original sample size, and the critical values correspond to .01 significance level. These critical values ('c Lévy' lines in the

[5]The row labelled χ^2 reproduces the asymptotic critical values for time series with finite moments.

[6]In the stable power GARCH model we would need the autocorrelation function of the unknown δ^{th} power of the time series. As $\delta < \alpha$ is a necessary condition for stationarity, we chose a value close to the upper bound of this region.

table), however, always depend on the 'unconditional' α, i.e., $\hat{\alpha}$ estimated from the raw data. This way we reproduce the 'model identification' phase. However, $\hat{\alpha}$ will in general be biased, if there is a significant autocorrelation in the return series, or in its powers. Assume, for example, that returns were simulated by a stable GARCH process with index α. In this case the estimated value of α—as a result of ignoring the GARCH dependence structure in the data—can, for example, be 0.2 smaller than that of the proper data generating process. This difference (which seems to be quite typical) increases the critical value corresponding to the first order Ljung-Box statistics by roughly 2, therefore it is more likely that the null hypothesis of the data being independent will be accepted. Even taking the above uncertainty into account, most statistics reported in the table are well below critical values for returns, and well above critical values for squared returns. Exceptions are third and fourth order statistics calculated from 30 minute returns, and from 60 minute squared returns. The latter result is irrelevant, and from the former one we might anticipate the presence of higher order ARMA terms, but that was later not confirmed by specification tests.[7]

Critical values reported in the 'c normal' rows of the tables correspond to the normal distribution. For comparison with the Lévy values, these were also obtained by Monte Carlo simulations, but they are close to the asymptotic values. Lévy critical values are typically greater than the normal ones, and the difference is more marked for critical values calculated from squared returns.

In light of the above results for CISCO we took the ARMA part of the stable GARCH model (μ_t) to be constant. We did not attempt to identify the orders of GARCH because we have no knowledge of the distribution of the sample autocorrelations of the δ^{th} power of returns ($\hat{\rho}_{n,X^\delta}(h)$), but even if we had, we would not know the value of delta prior to estimation. The results of Mikosch and Stărică (2000), concerning the distribution of $\hat{\rho}_{n,X^2}(h)$ for X being a GARCH(1,1) process with finite variance innovations are rather discouraging: they found that $\hat{\rho}_{n,X^2}(h)$ does not converge to a constant limit for a nearly integrated GARCH(1,1) process, so this statistics does not estimate anything.

We used 15, 30 and 60 minute returns of CISCO, and 15 minute returns of MOL in the year 1998 (full year) for estimation of the stable GARCH model. We only used 15 minute returns for MOL because the sample size for longer time

[7]It is well-known that an ignored GARCH process biases tests for residual autocorrelation towards rejection of no serial correlation in the residuals.

horizons was relatively small.[8] We also estimated two kinds of restrictions of the general model: normal power GARCH ($\alpha = 2$) and normal GARCH ($\alpha = \delta = 2$). We found that ARMA(0,0) GARCH(1,1) models best fit data in all cases (we used the AIC criterion for model selection).

Tables 4-6 (CISCO) and Table 8 (MOL) report estimated parameters, standard errors, upper and lower limits of 99% confidence intervals, persistence (V) and optimal value of the log likelihood function ($loglik$), and the Jarque–Bera test for normality of the innovations (JB-norm). Reported standard errors and confidence intervals were calculated by bootstrap.[9]

Looking at the results for CISCO we note that the estimated values of α slightly increase with increasing time horizons. The 60 minute α is outside of the 99% confidence interval for the 15 minute α; however, the 15 minute α is just on the lower limit of the 60 minute confidence interval. The Lévy GARCH model fits better than the normal to all data series, the $\alpha = 2$ restriction is always rejected at any meaningful significance level, and innovations are clearly not normally distributed. The value of α for MOL is lower than for CISCO, indicating that ouliers are more frequent among the Hungarian stock's returns. A possible reason for this might be the Russian crisis during the investigated period. Another reason may be the size of the market: a smaller market may well be more sensitive to shocks.

Estimated values of δ for CISCO seem to increase with longer time horizons. Delta is significantly less than alpha in all cases (a condition for strict stationarity). The distribution of the estimate of delta is skewed. An interesting phenomenon is that normal power GARCH coincides with normal GARCH for the 30 and 60 minute returns of CISCO. The value of persistence is one in Lévy power GARCH models, and slightly lower in other models. A high value of persistence means that the effect of shocks die out slowly on the market.

As we used relatively large frequency time series, high values of persistence are by no means surprising. However, it is interesting to note that Lévy GARCH always gives larger persistence values than the corresponding Gaussian GARCH models. It is also intriguing that Gaussian GARCH for CISCO does not give

[8] The Budapest Stock Exchange only operated for two hours daily in 1998. Opening times were extended in 1999.

[9] The number of repetitions was 1000, and simulated sample sizes agreed with the size of the data sample the model was estimated from.

decreasing persistence values for aggregated time series.

Drost and Nijman (1993) analysed the effect of temporal aggregation on GARCH models when (the i.i.d.) innovations have finite kurtosis. They proved that weak GARCH(1,1) is closed to temporal aggregation, i.e., the aggregated series will also follow a GARCH(1,1) process, and its persistence will be the k^{th} power of the persistence of the high frequency model, where k is the number of cumulated periods in the low frequency time series, provided that the persistence is less than one. Our Gaussian GARCH models satisfy these conditions, but the estimated persistencies do not follow this property.[10] This clearly indicates that the assumption of Gaussian innovations should be rejected.

Even though the derivation in Drost and Nijman (1993) does not apply to GARCH models with Lévy distributed innovations, as the higher moments of the distribution do not exist unless $\alpha = 2$ (i.e., the Gaussian case), the intuition behind their theorem strongly points to the case that stable GARCH models should also be closed to aggregation. And that should also mean that α should be constant; however, our results for CISCO do not contradict to that conclusion.

[10]Given that persistence is roughly 0.9 in the 15 minute model, it should be .81 in the 30 minute model, and .66 in the 60 minute model. However, our estimates for the 30 and 60 minutes persistence are higher than that of the 15 minutes series, c.f., Table 6.

Table 1: Response surface of the Ljung–Box test statistic to α, original series: (z_t)

α	$p=1$		$p=2$		$p=3$		$p=4$	
	0.95	0.99	0.95	0.99	0.95	0.99	0.95	0.99
χ^2	3.84	6.64	5.99	9.21	7.82	11.34	9.49	13.28
2	3.82	6.56	5.95	9.23	7.83	11.35	9.53	13.28
1.95	3.79	6.57	5.91	9.11	7.71	11.34	9.40	13.16
1.9	3.75	6.53	5.85	9.19	7.64	11.34	9.30	13.38
1.85	3.62	6.55	5.78	9.26	7.60	11.55	9.24	13.57
1.8	3.53	6.46	5.61	9.36	7.37	11.83	9.04	13.84
1.75	3.38	6.33	5.42	9.38	7.22	12.04	8.86	14.53
1.7	3.20	6.52	5.31	9.92	7.08	13.00	8.75	15.74
1.65	3.10	6.77	5.12	10.66	6.97	14.45	8.67	18.22
1.6	2.89	6.55	4.90	11.03	6.69	15.37	8.36	19.09
1.55	2.66	6.72	4.64	11.81	6.46	17.31	8.20	22.00
1.5	2.43	6.88	4.33	12.76	6.13	19.16	7.93	24.58
1.45	2.18	6.84	4.02	13.41	5.84	19.93	7.65	26.18
1.4	1.96	6.78	3.69	14.58	5.47	21.66	7.20	27.92
1.35	1.70	6.64	3.31	14.07	5.05	22.05	6.80	30.06
1.3	1.44	6.55	2.89	13.78	4.47	21.61	6.15	29.78
1.25	1.23	6.55	2.59	15.02	4.10	23.26	5.77	31.85
1.2	0.99	6.02	2.18	13.56	3.58	21.94	5.04	30.62
1.15	0.77	4.69	1.76	12.54	2.97	21.05	4.39	30.39
1.1	0.60	4.54	1.43	12.14	2.56	20.49	3.87	30.19

χ^2 gives the asymptotic critical values for distributions with finite moments. We used 2000 observations for the time series. The number of repetitions was 100000.

Table 2: Response surface of the Ljung–Box test statistic to α, power series: $(|z|_t^\delta)$

α	$p=1$		$p=2$		$p=3$		$p=4$	
	0.95	0.99	0.95	0.99	0.95	0.99	0.95	0.99
χ^2	3.84	6.64	5.99	9.21	7.82	11.34	9.49	13.28
2	3.77	6.64	5.94	9.38	7.77	11.49	9.45	13.41
1.95	2.46	6.24	4.36	10.41	6.13	14.52	7.91	18.33
1.9	1.68	6.43	3.37	12.83	5.08	18.77	6.80	24.74
1.85	1.28	6.50	2.80	14.46	4.48	23.38	6.23	30.54
1.8	1.06	6.60	2.50	15.58	4.16	25.19	6.04	33.17
1.75	0.93	6.80	2.28	15.84	3.98	25.33	5.80	34.90
1.7	0.85	7.15	2.20	16.39	3.95	26.09	5.84	35.02
1.65	0.79	7.03	2.10	16.40	3.83	27.12	5.82	36.58
1.6	0.74	7.07	2.02	16.47	3.69	26.26	5.62	33.72
1.55	0.74	6.77	2.01	16.37	3.68	26.43	5.63	37.43
1.5	0.71	6.75	1.96	16.37	3.61	26.90	5.55	36.04
1.45	0.70	7.15	1.99	17.63	3.70	26.91	5.62	36.84
1.4	0.67	6.73	1.94	16.66	3.63	27.42	5.62	36.68
1.35	0.67	7.38	1.96	17.59	3.70	28.20	5.77	37.58
1.3	0.67	6.84	1.94	17.28	3.67	27.39	5.50	37.32
1.25	0.65	7.25	1.87	16.69	3.66	26.10	5.58	35.64
1.2	0.65	7.24	1.90	17.79	3.61	27.48	5.64	39.47
1.15	0.64	7.08	1.85	17.57	3.61	28.00	5.69	39.16
1.1	0.62	6.70	1.81	17.25	3.49	27.12	5.42	37.04

χ^2 gives the asymptotic critical values for distributions with finite moments. We used 2000 observations for the time series. The number of repetitions was 100000. The value of δ was set to $\alpha - .001$ in all cases.

Table 3: Response surface of the Ljung–Box test statistic to α, squared series: (z_t^2)

α	$p=1$		$p=2$		$p=3$		$p=4$	
	0.95	0.99	0.95	0.99	0.95	0.99	0.95	0.99
χ^2	3.84	6.64	5.99	9.21	7.82	11.34	9.49	13.28
2	3.77	6.64	5.94	9.38	7.77	11.49	9.45	13.41
1.95	2.38	6.21	4.27	10.54	6.05	14.77	7.80	18.70
1.9	1.50	6.19	3.09	12.75	4.75	19.00	6.40	24.97
1.85	1.02	5.93	2.33	14.39	3.87	23.25	5.49	30.81
1.8	0.74	5.63	1.86	14.49	3.22	24.27	4.94	32.62
1.75	0.56	5.48	1.50	13.84	2.83	23.83	4.34	33.72
1.7	0.46	5.37	1.28	13.96	2.51	24.00	4.00	32.84
1.65	0.36	4.68	1.08	13.38	2.18	24.35	3.58	34.36
1.6	0.30	4.38	0.89	13.10	1.87	21.85	3.09	29.93
1.55	0.25	3.87	0.77	11.42	1.65	21.39	2.72	33.67
1.5	0.21	3.15	0.64	11.01	1.39	20.38	2.43	30.54
1.45	0.18	3.15	0.55	11.21	1.20	20.07	2.13	28.84
1.4	0.15	2.54	0.46	9.37	0.99	18.34	1.87	27.90
1.35	0.13	2.57	0.38	9.86	0.86	17.73	1.61	26.60
1.3	0.11	2.06	0.31	8.36	0.70	16.60	1.26	25.73
1.25	0.09	1.82	0.25	7.04	0.55	13.92	1.08	22.40
1.2	0.07	1.56	0.20	6.35	0.44	13.85	0.88	23.23
1.15	0.06	1.22	0.17	5.79	0.37	13.17	0.70	21.10
1.1	0.05	0.78	0.13	4.55	0.28	10.39	0.53	17.32

χ^2 gives the asymptotic critical values for distributions with finite moments. We used 2000 observations for the time series. The number of repetitions was 100000.

Table 4: Estimated parameters of the Lévy power GARCH model, CISCO

Frequency	Nob	$10^6\mu$	α	δ	$10^6 c_0$	c_1	d_1	V	*loglik*
15 min	6548	69	1.75	1.50	14	0.107	0.721	1.000	26719
		45	0.02	0.06	10	0.008	0.016		
		182	1.80	1.58	72	0.130	0.760		
		-58	1.71	1.22	9	0.088	0.681		
30 min	3274	91	1.77	1.54	6	0.036	0.905	1.000	12137
		99	0.02	0.07	4	0.005	0.009		
		347	1.84	1.68	29	0.049	0.926		
		-152	1.72	1.27	3	0.024	0.874		
60 min	1637	141	1.83	1.66	8	0.034	0.899	1.000	5448
		198	0.03	0.12	51	0.008	0.018		
		670	1.91	1.82	143	0.063	0.937		
		-343	1.75	1.15	3	0.018	0.841		

Estimated parameters appear in the first row of each frequency. Standard errors, upper and lower limits of 99% confidence intervals are reported in rows 2, 3 and 4, respectively.

Table 5: Estimated parameters of the normal power GARCH model, CISCO

Frequency	Nob	$10^6\mu$	δ	$10^6 c_0$	c_1	d_1	V	*loglik*	JB-norm
15 min	6548	60	1.69	7	0.129	0.685	0.894	26496	1864
		50	0.16	13	0.011	0.020			
		174	2.00	76	0.155	0.731			
		-69	1.28	1	0.103	0.632			
30 min	3274	70	2.00	1	0.035	0.901	0.971	12026	720
		109	0.09	1	0.005	0.014			
		343	2.00	10	0.052	0.932			
		-213	1.47	0	0.024	0.861			
60 min	1637	72	2.00	2	0.044	0.873	0.960	5409	281
		218	0.15	10	0.010	0.028			
		588	2.00	70	0.076	0.920			
		-453	1.32	1	0.025	0.780			

Estimated parameters appear in the first row of each frequency. Standard errors, upper and lower limits of 99% confidence intervals are reported in rows 2, 3 and 4, respectively.

Table 6: Estimated parameters of the normal GARCH model, CISCO

Frequency	Nob	$10^6\mu$	$10^6 c_0$	c_1	d_1	V	*loglik*	JB-norm
15 min	6548	61	1	0.113	0.671	0.898	26494	1906
		49	0	0.007	0.020			
		173	2	0.133	0.718			
		-65	1	0.094	0.617			
30 min	3274	70	1	0.035	0.901	0.971	12026	720
		109	0	0.005	0.013			
		331	1	0.047	0.930			
		-229	0	0.022	0.862			
60 min	1637	71	2	0.044	0.873	0.960	5409	281
		216	1	0.008	0.022			
		627	4	0.065	0.923			
		-483	1	0.024	0.800			

Estimated parameters appear in the first row of each frequency. Standard errors, upper and lower limits of 99% confidence intervals are reported in rows 2, 3 and 4, respectively.

Table 7: Ljung-Box test statistics and critical values, CISCO

	15 min				30 min				60 min			
n	1	2	3	4	1	2	3	4	1	2	3	4
y_t	0.16	0.85	0.93	0.93	0.44	3.20	15.3	21.5	0.04	0.49	0.55	1.25
c (Lévy)	6.60	10.1	14.6	18.2	7.78	10.6	13.7	16.5	6.57	11.4	14.8	16.7
c (normal)	6.53	8.94	10.8	13.0	6.64	8.84	11.7	12.6	7.56	9.47	11.5	12.6
y_t^2	501	781	994	1208	59.3	92.6	419	426	12.8	25.4	27.0	32.6
c (Lévy)	2.65	17.2	20.6	47.7	2.35	14.0	19.0	34.0	6.16	23.3	33.2	39.5
c (normal)	6.87	8.77	11.1	13.0	6.28	9.08	11.3	12.5	6.85	8.32	11.0	12.6
Lévy power GARCH residuals												
$\|z_t\|^\delta$	0.24	0.29	0.41	0.51	6.09	29.0	35.5	42.0	7.78	8.73	14.4	14.4
c	5.56	18.6	29.3	47.5	4.74	15.8	21.6	37.4	7.35	23.1	30.8	37.8
z_t^2	0.43	0.43	1.01	1.18	2.88	46.2	52.4	55.6	8.25	8.83	11.7	11.8
c	2.65	17.2	20.6	47.7	2.35	14.0	19.0	34.0	6.16	23.3	33.2	39.5
Normal power GARCH residuals												
$\|z_t\|^\delta$	0.05	0.16	0.19	0.20	1.41	9.60	17.3	22.5	1.42	2.97	10.2	10.6
c	7.21	9.05	11.2	13.3	6.28	9.08	11.3	12.5	6.85	8.32	11.0	12.6
z_t^2	0.07	0.10	0.29	0.34	1.41	9.60	17.3	22.5	1.42	2.97	10.2	10.6
c	6.87	8.77	11.1	13.0	6.28	9.08	11.3	12.5	6.85	8.32	11.0	12.6
Normal GARCH residuals												
z_t^2	0.50	0.51	0.72	0.87	1.41	9.59	17.3	22.5	1.42	2.97	10.2	10.6
c	6.87	8.77	11.1	13.0	6.28	9.08	11.3	12.5	6.85	8.32	11.0	12.6

n denotes the order of statistics. Statistics calculated from various powers of returns (y_t) and residuals (z_t) appear in the rows y_t, y_t^2, $|z_t|^\delta$, and z_t^2. Below the statistics critical values (c) corresponding to 99% probability appear (see text).

Table 8: Parameters of Lévy power GARCH and normal GARCH models estimated from 15 minute returns of MOL, as well as Ljung-Box test statistics and critical values calculated from powers of model residuals.

	$10^6\mu$	α	δ	10^6c_0	c_1	d_1	V	$loglik$	JB-norm
Lévy	129	1.38	0.85	125	0.054	0.911	0.998	6426.3	10862
	61	0.03	0.07	87	0.008	0.009			
	297	1.47	1.03	583	0.077	0.936			
	-3	1.31	0.64	45	0.034	0.885			
z_t		0.00 (9.41)		0.03 (32.7)		0.05 (35.9)		0.07 (36.8)	
$\|z_t\|^\delta$		0.02 (11.5)		0.03 (29.7)		0.03 (33.6)		0.03 (38.1)	
z_t^2		0.00 (3.69)		0.00 (19.7)		0.00 (21.0)		0.00 (21.7)	
normal	1574			5	0.039	0.805	0.883	5947.0	6539
	214			2	0.010	0.063			
	2109			14	0.065	0.911			
	1028			2	0.014	0.559			
z_t		2.80 (6.18)		2.81 (8.43)		2.95 (11.1)		3.08 (12.2)	
z_t^2		0.79 (6.35)		2.01 (8.46)		2.80 (10.4)		2.97 (12.6)	

Estimated parameters appear in the rows labelled 'Lévy' and 'Normal'. Standard errors, upper and lower limits of 99% confidence intervals are reported in rows 2, 3 and 4, respectively. Ljung-Box statistics (up to fourth order) calculated from powers of model residuals (z_t) appear in rows $|z_t|^\delta$ and z_t^2, followed by critical values (in parentheses) corresponding to 99% probability level. Sample size is 1798.

We calculated Ljung-Box test statistics from powers of model residuals (z_t, $|z_t|^\delta$ and z_t^2) to check if there was significant autocorrelation left in the filtered data. (δ was set to the estimated value.) Results are reported in Tables 7-8. As before, critical values were calculated from Monte Carlo simulations. For 15 minute return residuals of CISCO (Table 7.) statistics are well below critical values. For the Lévy power GARCH residuals of 60 minute returns of CISCO, however, the first order statistics are slightly higher than the critical values. Interestingly, the normal GARCH model does better in this case, but neither model does well on the 30 minute returns. We do not report the test statistics calculated for residual autocorrelation of the CISCO returns; these were almost zero: well below critical values for all frequencies. For the stock MOL (15 minute data, Table 8.) all statistics were well below critical values; the Lévy model seems to do slightly better in this respect.

Finally, we checked if residuals of the Lévy power GARCH models fitted to 15 minute CISCO and MOL returns were stable under addition. (We used residuals from 15 minute returns partly because these constitute the largest sample, and partly because these (and their powers) seemed to contain the least amount of autocorrelation (see Ljung-Box statistics). This point is cruical - as we saw in the example in the introduction, data containing serial autocorrelation in the powers may appear to be stable under addition even if they do not come from a stable process. We calculated Monte Carlo confidence intervals for α estimated from aggregated residuals at the first 6 orders of aggregation. (We simulated 1000 samples of iid α stable random numbers with α set to the value estimated from the residuals, and at each level of aggregation we calculated the average, the upper and lower percentiles of α's. Simulated samples had the same size as the samples of residuals.)

Results are reported in Table 9. n is the level of aggregation, rows labeled CISCO and MOL contain α's estimated from aggregated residuals. Below these values the average, the upper and lower confidence limits for the Lévy distributed time series are reported.

α indices estimated from aggregated residuals converge to two quickly, and all (except MOL, $n = 2$) lie outside respective confidence intervals. In contrast to that, the average of α's estimated from aggregates of simulated stable random numbers remain practically constant. The width of confidence intervals increases

Table 9: Monte Carlo test of stability of residuals of the Lévy power GARCH model

n	1	2	3	4	5	6
CISCO	1.75	1.89	1.94	1.96	1.98	2.00
	1.75	1.75	1.75	1.75	1.75	1.75
	1.79	1.81	1.82	1.84	1.85	1.87
	1.70	1.68	1.66	1.65	1.65	1.64
MOL	1.39	1.43	1.58	1.75	1.91	1.95
	1.39	1.39	1.39	1.39	1.39	1.40
	1.49	1.51	1.55	1.58	1.61	1.63
	1.29	1.26	1.24	1.21	1.20	1.18

with the level of aggregation, as sample sizes decrease.

4 Conclusion

With efficiently computerised trading systems becoming standard even on small markets, and with cheaply and abundantly available information, there are increasingly more agents (in particular, day traders) on exchange markets who are interested in the short-run risk of their investments. Extreme events may matter more for day-traders than for strategic investors, interested in long-run returns. However, GARCH models based on normally distributed innovations may give less reliable prediction of the expected risk than models assuming stable distributed innovations.

One difficulty when developing a Lévy power GARCH model is that the standard identification and diagnostic tools cannot be applied the same way as with more regular GARCH models. This means that users of stable GARCH models will have to compute statistics by bootsrap. It seems that the estimated α value depends on the GARCH model itself, so the traditional identification procedure clearly cannot be used.

In our study we found that even though Lévy GARCH models fit to returns better than normal GARCH does, we can reject stability of the residuals. On the other hand Mittnik, Paolella and Rachev (2000), for example, reported results on a return series, for which stability of the model residuals could not be rejected. Both results are rather particular, and further studies are needed before a generic conclusion could be drawn about the stability of stable GARCH model residuals.

Even though our study demonstrated that the stability of residuals can be rejected in several cases, GARCH models with Lévy innovations still provide a better representation of the underlying data generating process than those driven by Gaussian distributions. Apparently, Lévy models are vulnerable to stability tests.[11] But all models are just specific approximations of the true data generating process, and stable GARCH models proved clearly superior to the Gaussian ones in all cases.

Besides examining the goodness-of-fit of models and stability of residuals one should compare models' performance in financial applications like option pricing, or value at risk calculations. A question one may consider for example is this: how much would the result of a value at risk calculation based on a Lévy GARCH model (say with $\alpha = 1.7$) be affected by the fact that residuals are not stable?

Another very important result is that the gain in using a Lévy GARCH model is much larger for MOL. Thus, extreme events are more likely to drive an emerging market than a mature one. Further, Gaussian GARCH estimated from MOL is much more sensitive to sample adjustments than from CISCO, which also indicates that risk analysis, based on Gaussian innovations, may be very misleading on an emerging capital market.

[11] Empirical papers using Gaussian GARCH models provide ample evidence against the normality of the innovation process.

References

[1] Aït Sahalia, Y., Y. Wang and F. Yared, 2001, Do Option Markets Correctly Price the Probabilities of Movements of the Underlying Asset?; Journal of Econometrics, Vol. 102, pp. 67-110.

[2] Bollerslev, T., 1986, Generalized Autoregressive Conditional Heteroskedasticity; Journal of Econometrics, Vol. 31, pp. 307–327.

[3] Bollerslev, T., R. F. Engle and D. B. Nelson, 1994, Arch Models; in: Handbook of Econometrics, Eds: Engle, R. F. and D. L. McFadden, Vol. 4, Chapter 49, North-Holland.

[4] Campbell, J. Y., Lo, A. W. and MacKinlay, A. C., 1997, The Econometrics of Financial Markets; Princeton University, New Jersey.

[5] Ding, Z., C. W. J. Granger and R. F. Engle, 1993, A Long Memory Property of Stock Market Returns and a New Model; Journal of Empirical Finance, Vol. 1, pp. 83-106.

[6] Drost, F. C. and T. E. Nijman, 1993, Temporal aggregation of GARCH processes; Econometrica, Vol. 61, No. 4, pp. 909–927.

[7] Ghose, D. and K. F. Kroner, 1995, The relationship between GARCH and symmetric stable processes: finding the source of fat tails in financial data, Journal of Empirical Finance, Vol. 2, pp. 225–251.

[8] Gourieroux, C., 1997, ARCH Models and Financial Applications; Springer.

[9] Johnson, S. and A. Schleifer, 1999, Coase vs. the Coasians; Harvard U research paper No. 1885.

[10] LePage, R., Podgórski, K. and Ryznar, M., 2001, Bootstrapping signs and permutations for regression with heavy tailed errors: a seamless resampling; mimeo.

[11] Liu, S. and Brorsen, B. W., 1995, Maximum Likelihood Estimation of a GARCH-Stable Model; Journal of Applied Econometrics, Vol. 10, pp. 273–285.

[12] McCulloch, J. H. , 1985, Interest-Risk Sensitive Deposit Insurance Premia: Stable ACH Estimates, Journal of Banking and Finance, Vol. 9, pp. 137–156.

[13] McCulloch, J. H. , 1998, Numerical Approximation of the Symmetric Stable Distribution and Density; In: Adler, R. J., Feldman, R. E. and Taqqu, M. S. eds. A Practical Guide to Heavy Tails, Birkhäuser, Boston, pp. 489–499.

[14] Mikosch, T. and Stărică, C., 2000, Limit Theory for the Sample Autocorrelations and Extremes of a GARCH(1,1) Process, The Annals of Statistics, Vol. 28 (5), pp. 1427–1451.

[15] Mittnik, S., Paolella, M. S. and Rachev, S. T., 2000, Diagnosing and Treating the Fat Tails in Financial Returns Data; Journal of Empirical Finance, Vol. 7, pp. 389–416.

[16] Mittnik, S., Paolella, M. S. and Rachev, S. T., 2002, Stationarity of Stable Power-GARCH Processes, Journal of Econometrics, Vol. 106, pp. 97–107.

[17] Palágyi, Z. and Mantegna, R. N. , 1999, Empirical Investigation of Stock Price Dynamics in an Emerging Market; Physica A, Vol. 269, pp. 132–139.

[18] Palágyi, Z., G. Kőrösi and R. N. Mantegna, 2001, High Frequency Data Analysis in an Emerging and a Developed Market; in: Takayasu, H., Ed., *Empirical Science of Financial Fluctuations*; Springer-Verlag, Tokyo.

[19] Panorska, A., Mittnik, S. and Rachev, S. T., 1995, Stable GARCH Models for Finacial Time Series; Applied Mathematics Letters, Vol 8, pp. 33–37.

[20] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992): Numerical Recepies in FORTRAN: The Art of Scientific Computing; 2nd ed. Cambridge University Press.

[21] Rachev, S. and S. Mittnik, 1999, Stable Paretian Models in Finance; Wiley.

[22] Samorodnitsky, G. and Taqqu, M. S., 1994, Stable non-Gaussian random processes, Chapman & Hall, New York, London.