Time consistency of Lévy models

Ernst Eberlein and Fehmi Özkan

University of Freiburg

Nr. 79 September 2002

rev. Nov. 2002

Department of Mathematical Stochastics
University of Freiburg
Eckerstraße 1
D–79104 Freiburg im Breisgau

and

Freiburg Center for Data Analysis and Modeling
University of Freiburg
Eckerstraße 1
D–79104 Freiburg im Breisgau

eberlein@stochastik.uni-freiburg.de

Time consistency of Lévy models

Ernst Eberlein
Fehmi Özkan

Department of Mathematical Stochastics, and Freiburg Center for Data Analysis and Modeling (FDM), University of Freiburg, Eckerstr. 1, D-79104 Freiburg, Germany

Key words: Lévy processes, scaling, intraday modeling, convolution, infinitely divisible distributions, generalized hyperbolic distributions

1 Introduction

Statistical accuracy of the models used in finance is vital for many areas such as pricing of derivatives, risk management and portfolio optimization. In risk management improving accuracy will result in a reduction of the number and and the severity of losses as well as in a potential freeing-up of economic and regularity capital. In portfolio optimization the ability to measure and control risk as well as chances with higher precision will lead to a more efficient allocation of the funds. The importance of accurate derivative pricing is evident. In order to improve the distributional properties of models for stock prices or indices many sophisticated generalizations of the classical geometric Brownian motion

$$\mathrm{d}S_t = \mu S_t \ \mathrm{d}t + \sigma S_t \ \mathrm{d}B_t \tag{1.1}$$

are used by practitioners. Within the diffusion world, μS_t and σS_t are replaced by more general processes $\mu(t, S_t)$ and $\sigma(t, S_t)$, but the driving Brownian motion $B = (B_t)_{t\geq 0}$ is retained. One of the key properties of Brownian motion is its selfsimilarity, i.e. if for any $c > 0$, the finite-dimensional distributions of $(B_{ct})_{t\geq 0}$ are identical to the finite-dimensional distributions of $\left(c^{1/2}B_t\right)_{t\geq 0}$. This beautiful scaling property is highly undesirable in finance. In the case of model (1.1) the distribution of the log returns $\log S_{t+\Delta} - \log S_t$ is normal, independent of the time span Δ considered. Only the two parameters of the distribution change in a linear way. Yet as is well known from a number of studies (see e.g. Fig. 7 in Eberlein and Keller (1995)) and as we will show later, the shape of empirical distributions varies considerably as a function of Δ. For widely spaced data such as security prices sampled on a monthly time grid, the return distribution is close to normal. Daily return distributions typically are far from normality. They have more mass at the origin and in the tails. The mass is taken away from the flanks. The distance to the normal distribution increases further if one goes down to an intraday time grid. Let us emphasize that the deviation from normality is not restricted to returns from equity data. It can be observed for returns from bond data as well (Raible (2000)). Note that for diffusions with general coefficients $\mu(t, S_t)$ and $\sigma(t, S_t)$ the distribution of S_t or of the derived log returns are even not known. Thus to test if the model is realistic, in these cases one is dependent on simulations.

If historical price data is used to choose the model parameters, the calibration is usually done on the basis of data which was collected on a fixed time grid, typically daily closing or daily Kassa prices. Once a model is fixed it is used to make predictions for time horizons different from the one underlying the calibration. Implicitly one assumes that the distributions produced by the model at these time horizons coincide with those one would get from the corresponding data. A typical example here is the estimation of the Value at Risk (VaR) from a daily return distribution. This value is then transformed with the square-root-of-ten-rule in a VaR-number for a horizon of ten trading days which corresponds to two weeks. It is the purpose of this paper to elaborate on this question. We shall discuss only models for equity prices and indices. The corresponding interest rate and credit risk models are developed in Eberlein and Raible (1999) and Eberlein and Özkan (2003). It will become clear that models driven by Lévy processes respond in the proper way to changes of time horizons, they are "time consistent". As a consequence this type of model produces more accurate estimates and prices. As one application it allows an accurate derivation of intraday risk. The question of time consistency of the model is also crucial if one wants to integrate market and credit risk since both are measured on totally different time scales.

2 The Lévy model

A Lévy process $X = (X_t)_{t\geq 0}$ is a process with stationary and independent increments. Underlying is a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ to which the process is adapted. The distribution of a Lévy process is completely determined by any of its one-dimensional marginal distributions P^{X_t}, say P^{X_1}. From the independence and the stationarity of the increments $X_t - X_s$ it follows that P^{X_1} is infinitely divisible, i.e. for any $n > 0$, P^{X_1} can be written as the n-fold convolution product $\nu_n * \cdots * \nu_n$. Conversely any infinitely divisible distribution ν generates in a natural way a Lévy process $X = (X_t)_{t\geq 0}$ which is characterized by the fact that $\nu = P^{X_1}$. Thus there is a one-to-one relationship between infinitely divisible probability measures and Lévy processes. Note that if such a probability measure is given by a set of parameters, the same holds for the corresponding Lévy process. The most prominent example is the Brownian motion which is characterized by the two parameters μ and σ^2 of the normal distribution $N(\mu, \sigma^2)$. In general a Lévy process can be decomposed in the following way

$$X_t = \alpha t + \sigma B_t + Z_t + \sum_{s\leq t} \triangle X_s 1\!\!1_{\{|\triangle X_s|\geq 1\}} \tag{2.1}$$

where $(B_t)_{t\geq 0}$ represents a standard Brownian motion, $(Z_t)_{t\geq 0}$ a purely discontinuous martingale, which is independent of $(B_t)_{t\geq 0}$ and $\triangle X_s := X_s - X_{s-}$ denotes the jump at time s. (2.1) is the so-called canonical semimartingale representation. It shows important structural properties of Lévy processes, namely the decomposition in a martingale – as opposed to a local martingale in general – and the bounded variation part, which consists of a linear drift αt and a sum of isolated jumps larger than a constant which is chosen to be 1 here. We shall focus later on Lévy processes where the generating infinitely divisible distribution has finite moments of all orders. In this case the last summand, which represents the sum of the big jumps can be merged with the drift term αt to a drift $\mathbb{E}[X_1]t$, i.e. α is then the mean of the generating distribution. Furthermore if in addition $\sigma = 0$, i.e. if the continuous Brownian part disappears, we have a process of the form $X_t = \mathbb{E}[X_1]t + Z_t$. Making use of the random measure of jumps of a process $(X_t)_{t\geq 0}$ defined by

$$\mu^X(\omega; \mathrm{d}t, \mathrm{d}x) = \sum_{s>0} 1\!\!1_{\{\triangle X_s(\omega)\neq 0\}}\, \varepsilon_{(s,\triangle X_s(\omega))}(\mathrm{d}t, \mathrm{d}x)$$

where $\varepsilon_{(s,\triangle X_s(\omega))}$ denotes the Dirac measure at $(s, \triangle X_s(\omega))$, we get the following representation

$$X_t = \mathbb{E}[X_1]t + \int_0^t \int_{\mathbb{R}\setminus\{0\}} x \left(\mu^X(.; \mathrm{d}u, \mathrm{d}x) - \mathrm{d}uF(\mathrm{d}x)\right). \tag{2.2}$$

Here F denotes the Lévy measure of the generating distribution $\nu = P^{X_1}$. The stochastic integral is well defined, since X is a special semimartingale (Jacod and Shiryaev (1987) or Skorohod (1991)).

The property we want to see primarily in a model for equity prices is that the distribution of log returns produced by the model along a time interval of length 1 is exactly the one which we get from fitting empirical log returns. Instead of length 1 of course any other fixed length Δ could be considered as well. Denote by ν the fitted distribution. If ν is an infinitely divisible distribution – which represents a very large class – our goal is achieved by choosing the price process $S = (S_t)_{t\geq 0}$ in the form

$$S_t = S_0 \exp\left(X_t\right) \tag{2.3}$$

where $(X_t)_{t\geq 0}$ is the Lévy process generated by ν (Eberlein and Keller (1995)). The classical model (1.1) is of this type since equation (1.1) has the solution $S_t = S_0 \exp\left(\sigma B_t + \left(\mu - \sigma^2/2\right) t\right)$. Equation (2.3) can also be described by the stochastic differential equation

$$\mathrm{d}S_t = S_{t-}\left(\mathrm{d}X_t + \frac{1}{2}\sigma^2\ \mathrm{d}t + \int_{\mathbb{R}} \left(\mathrm{e}^x - 1 - x\right)\ \mu^X(\mathrm{d}t, \mathrm{d}x)\right). \tag{2.4}$$

Essentially any extension of (1.1) which has been discussed in the literature can be considered in this more general framework as well. We just mention the stochastic volatility models in Eberlein, Kallsen, and Kristen (2003).

It would be tempting to take the stochastic differential equation (1.1) as the starting point for a more general model and replace $(B_t)_{t\geq 0}$ by a Lévy process $(X_t)_{t\geq 0}$ in this equation. Because of the possible jumps of $(X_t)_{t\geq 0}$ one has then to replace S_t in (1.1) by the left limits S_{t-}. The solution of the equation $\mathrm{d}\widetilde{S}_t = \widetilde{\mu}\widetilde{S}_{t-}\ \mathrm{d}t + \widetilde{\sigma}\widetilde{S}_{t-}\ \mathrm{d}X_t$ is no longer an ordinary exponential, but the stochastic or Doléans-Dade exponential. Its explicit form can be deduced by using Itô's formula for semimartingales (see Jacod and Shiryaev (1987))

$$\widetilde{S}_t = \mathrm{e}^{(\widetilde{\mu} - (\widetilde{\sigma}\sigma)^2/2)t + \widetilde{\sigma} X_t} \prod_{s\leq t} (1 + \widetilde{\sigma}\triangle X_s)\mathrm{e}^{-\widetilde{\sigma}\triangle X_s}.$$

The distribution of the log returns along time length 1 does not coincide with the distribution generating $(X_t)_{t\geq 0}$ any more. $(\widetilde{S}_t)_{t\geq 0}$ could also have negative values, unless we consider only Lévy processes X satisfying $\widetilde{\sigma}\triangle X > -1$, which is highly undesirable. But still it is possible to follow this approach (see e.g. Chan (1999)). Note that the restriction concerning negative jumps is redundant in our modeling approach (2.3).

Theoretically one could take the full class of infinitely divisible distributions when fitting data. These are completely characterized by the triplet (b, c, F) which appears in the Lévy-Khintchine representation of the Fourier transform

$$\phi(u) = \exp\left[iub - \frac{1}{2}u^2 c + \int_{-\infty}^{\infty} \left(\mathrm{e}^{iux} - 1 - iuxh(x)\right)\ \mathrm{d}F(x)\right] \tag{2.5}$$

where $h(x) = \mathbb{1}_{[-1,1]}(x)$. b and $c \geq 0$ are real numbers, whereas F is a measure and therefore rather difficult to estimate. A large and very flexible subclass which can be described by five real parameters only is given by the generalized hyperbolic (GH) distributions. This class was introduced by Barndorff-Nielsen (1977) in the context of the so-called sand project. In a series of papers (Eberlein and Keller (1995), Barndorff-Nielsen (1998), Raible (2000), Eberlein and Prause (2002)) this class turned out to provide excellent fits to financial returns from equity prices, indices, FX-rates and bond prices. The density of a generalized hyperbolic distribution is given by

(2.6)

$$d_{GH(\lambda,\alpha,\beta,\delta,\mu)}(x) = a(\lambda,\alpha,\beta,\delta)\left(\delta^2+(x-\mu)^2\right)^{(\lambda-1/2)/2} \times \mathrm{K}_{\lambda-1/2}(\alpha\sqrt{\delta^2+(x-\mu)^2}\,)\exp\left(\beta(x-\mu)\right) \qquad (x\in\mathbb{R})$$

with norming constant

$$a(\lambda,\alpha,\beta,\delta) = \frac{(\alpha^2-\beta^2)^{\lambda/2}}{\sqrt{2\pi}\alpha^{\lambda-1/2}\delta^\lambda \mathrm{K}_\lambda(\delta\sqrt{\alpha^2-\beta^2}\,)}.$$

Here K_λ denotes the modified Bessel function of the third kind with index λ and the parameter space is given by

$$\delta>0, \quad 0\le|\beta|<\alpha \quad \text{ and } \lambda,\mu\in\mathbb{R}.$$

λ is a class, and α a shape parameter, β describes skewness, μ location and δ is a scaling parameter. There are various alternative parameterizations with useful properties. Quite a number of standard distributions are either special cases or limiting cases of GH-distributions. Among those are the normal and the Gamma as well as the reciprocal Gamma, the normal inverse Gaussian (NIG) (see e.g. Barndorff-Nielsen (1998)), the Variance Gamma distribution (Madan and Milne (1991)), the skewed Laplace, the Cauchy, the Student-t distribution, and the generalized inverse Gaussian (GIG) distributions. A discussion of these issues with all analytical details is given in Eberlein and v. Hammerstein (2002).

One fact on which we will dwell here is that a good statistical fit is still obtained if one samples on an intraday grid. The following graph shows hourly returns from Bayer stock price fluctuations and the fitted hyperbolic as well as the fitted normal density. The hyperbolic distribution used here is the subclass represented by $\lambda=1$. The second graph shows the same densities on a logarithmic scale. This transformation is necessary to make the behavior in the tails visible.

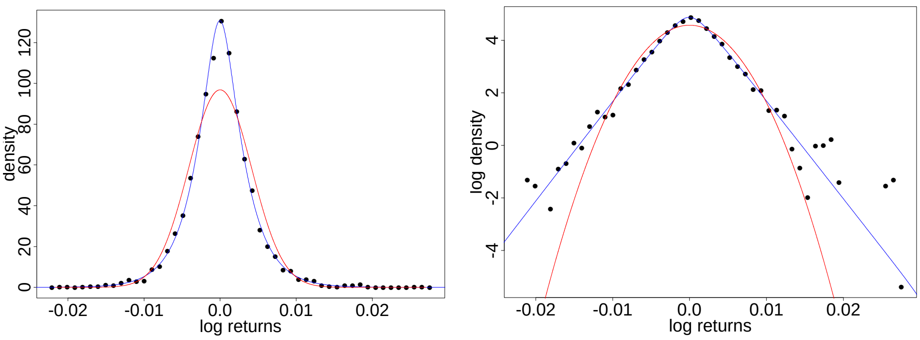

Figure 1: Empirical density of one-hour returns (Bayer) vs. density of fitted hyperbolic (blue) and fitted normal distribution (red).

The subclass of hyperbolic distributions is used here since due to the properties of the Bessel functions the density reduces to

(2.7)

$$d_{H(\alpha,\beta,\delta,\mu)}(x) = \frac{\sqrt{\alpha^2-\beta^2}}{2\alpha\delta\mathrm{K}_1\left(\delta\sqrt{\alpha^2-\beta^2}\right)}\exp\left(-\alpha\sqrt{\delta^2+(x-\mu)^2}+\beta(x-\mu)\right).$$

Table 1 shows the hyperbolic parameters obtained from fitting daily returns from various stocks by using maximum likelihood estimation. In order to simplify the programming we shall essentially focus on this subclass in the next sections. For estimation of parameters in the full class of GH distributions see Eberlein and Prause (2002).

	$\widehat{\alpha}$	$\widehat{\beta}$	$\widehat{\delta}$	$\widehat{\mu}$
Bayer	153.3	-6.2	0.0092	0.0012
Daimler	139.6	11.4	0.0123	-0.0019
Deutsche Bank	163.1	-7.6	0.0063	0.0007
Siemens	230.2	-4.1	0.0149	0.0005
Thyssen	107.7	-4.0	0.0110	0.0017

Table 1: Estimation of hyperbolic parameters for daily data.

An infinitely divisible distribution ν generates a convolution semigroup $(\nu_t)_{t\geq 0}$ of probability measures such that $\nu_1 = \nu$. These probability measures ν_t are the one-dimensional marginal distributions P^{X_t} of the Lévy process $X = (X_t)_{t\geq 0}$ connected to ν. The following graph shows the densities of ν_t as a function of t. Note that $\nu_0 = \varepsilon_0$, the Dirac measure with all mass concentrated at the origin. The convolution semigroup in the picture is generated by a hyperbolic distribution with density $d_{H(350,10,0.01,0)}$.

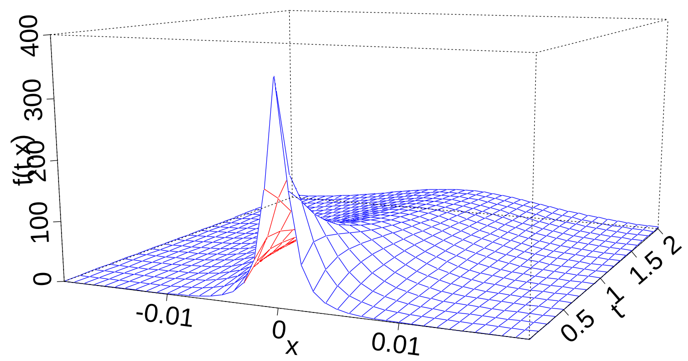

Figure 2: Convolution semigroup generated by the hyperbolic distribution with parameters $(\alpha, \beta, \delta, \mu) = (350, 10, 0.01, 0)$.

If ϕ denotes the Fourier transform of the generating distribution ν, then the Fourier transforms of ν_t are given by the t-th powers ϕ^t. Thus the densities of ν_t can be easily obtained via Fourier inversion. Let us also note that Fourier transforms $\phi(u) = \int \mathrm{e}^{iux} d_{GH}(x)\,\mathrm{d}x$ (as well as moment generating functions for arguments where they exist) can be very elegantly derived because of the exponential form of d_{GH}. One has just to merge e^{iux} with the density and to use that $\int d_{GH}(x)\,\mathrm{d}x = 1$

for any set of GH-parameters. The remaining expression is

$$\phi_{GH(\lambda,\alpha,\beta,\delta,\mu)} = \mathrm{e}^{iu\mu} \left(\frac{\alpha^2 - \beta^2}{\alpha^2 - (\beta + iu)^2} \right)^{\lambda/2} \frac{\mathrm{K}_\lambda(\delta\sqrt{\alpha^2 - (\beta + iu)^2})}{\mathrm{K}_\lambda\left(\delta\sqrt{\alpha^2 - \beta^2}\right)} \tag{2.8}$$

which is with the exception of the first factor the quotient of two normalizing constants for different choices of parameters. The moment generating function

$$M_{GH(\lambda,\alpha,\beta,\delta,\mu)}(u) = \mathrm{e}^{\mu u} \left(\frac{\alpha^2 - \beta^2}{\alpha^2 - (\beta + u)^2} \right)^{\lambda/2} \frac{\mathrm{K}_\lambda\left(\delta\sqrt{\alpha^2 - (\beta + u)^2}\right)}{\mathrm{K}_\lambda\left(\delta\sqrt{\alpha^2 - \beta^2}\right)} \tag{2.9}$$

exists for all u such that $|\beta + u| < \alpha$. Consequently generalized hyperbolic distributions have moments of all order. The class actually touches the borderline to distributions without finite moments. For example the Cauchy distribution appears as the limiting case for $\lambda = -0.5$ and $\alpha, \beta \to 0$. Because of the finiteness of the expectation $\mathbb{E}[GH]$ we can get rid of the truncation function $h(x)$ in the Lévy-Khintchine representation (2.5) and the drift coefficient becomes $\mathbb{E}[GH]$. Since $c = 0$ for GH-distributions the corresponding representation is

$$\phi_{GH}(u) = \exp\left[iu\mathbb{E}[GH] + \int_{-\infty}^{\infty} \left(\mathrm{e}^{iux} - 1 - iux\right) g_{GH}(x)\, \mathrm{d}x \right] \tag{2.10}$$

where g_{GH} denotes the Lévy density of the distribution. (2.10) completely determines the path properties of the generated Lévy process $(X_t)_{t\geq 0}$. Since $c = 0$ the process is purely discontinuous, i.e. the Brownian motion $(B_t)_{t\geq 0}$ in (2.1) disappears. The jump pattern of the paths is given by g_{GH}. If we write $\rho(x) = x^2 g_{GH}(x)$ then according to Raible (2000, Prop. 2.18) the behavior of $\rho(x)$ near the origin is given by

$$\rho(x) = \frac{\delta}{\pi} + \frac{\lambda + 1/2}{2}|x| + \frac{\delta\beta}{\pi}x + o(|x|) \qquad (x \to 0). \tag{2.11}$$

As an immediate application we see that the Lévy measure has infinite mass around the origin, i.e. $\int_{-1}^{+1} g_{GH}(x)\, \mathrm{d}x = \infty$. This means that the paths of the process have infinitely many jumps in any finite time interval. Since also $\int_{-1}^{+1} |x| g_{GH}(x)\, \mathrm{d}x = \infty$ the paths do not have finite variation. In the case of hyperbolic distributions used in Sections 3 and 4 one has to set $\lambda = 1$ in formulas (2.8) – (2.11). Figure 3 shows a simulation of a sample path of a GH-process. The picture resembles very much a high frequency intraday price path.

3 Fitting of intraday returns

The data set underlying the empirical investigation consists of complete intraday price records for a number of blue chips traded at the Frankfurt stock exchange. It covers the period from January 2, 1992 to August 19, 1994 and represents trading on the floor with opening at 10.30 a.m. and closing at 1.30 p.m. For each day all prices at which trades occurred are given together with the corresponding time stamp. From this data set we extracted log returns on various equidistant time grids. The statistical fitting is done via maximum likelihood estimation. Figure 1 shows an excellent fit of the hyperbolic density with parameters $(\alpha, \beta, \delta, \mu) = (377.74, 3.04, 0.0017, 0)$ to hourly Bayer log returns. The prices from which the hourly log returns are derived are those at time points 10.30 a.m. (opening), 11.30 a.m., 12.30 a.m., and 1.30 p.m. (closing). The QQ-plot on the left side

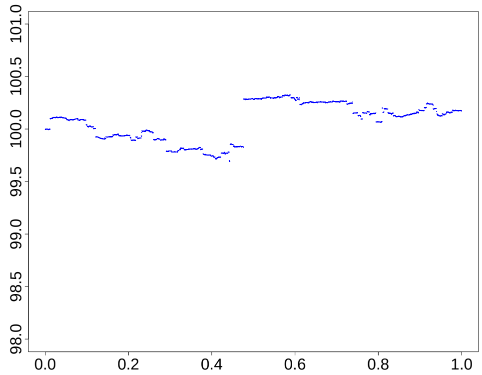

Figure 3: Simulation of a GH-path.

of Figure 4 compares one-hour log returns from Daimler data with the fitted hyperbolic distribution. The corresponding graph for 10-minutes log returns is given on the right side. Note that the small jump at the origin is the result of a certain number of 0-returns on such a narrow time grid. With more recent data sets taken from the electronic trading platform (XETRA) the number of trades goes up to several thousand per day and consequently 0-returns are rare events on a time grid of a few minutes. The remarkable fit which is obtained also for empirical distributions from narrow time grids, underlines the flexibility of the class of (generalized) hyperbolic distributions.

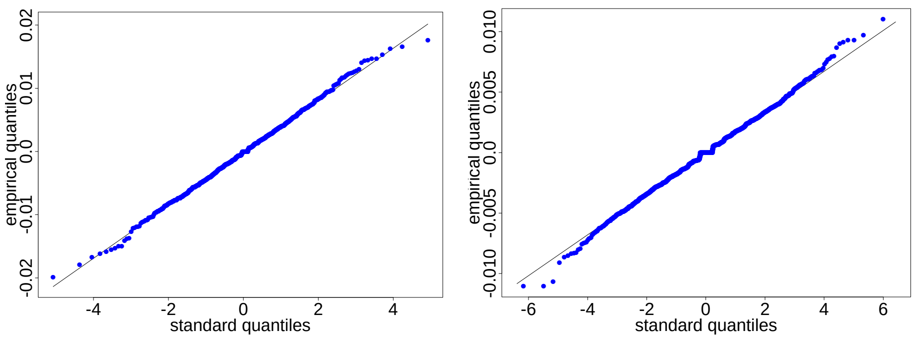

Figure 4: QQ-plot empirical density of one-hour log returns (Daimler) vs. fitted hyperbolic density (left) and QQ-plot of 10-minutes log returns (Daimler) vs. fitted hyperbolic distribution (right).

In addition to the graphical comparison we want to give some more quantitative results. First we show in Table 2 frequency distributions for stock prices of five companies. In the columns the relative frequencies of the returns in the interval $]-k\sigma, k\sigma[$ are compared with the probabilities of the fitted hyperbolic distributions $H(\widehat{\alpha}, \widehat{\beta}, \widehat{\delta}, \widehat{\mu})$. The tails starting at five standard deviations are given in the column labeled by $> 5\sigma$, where the numbers in the brackets give the absolute frequencies found in this region resp. the expected value for the corresponding hyperbolic distribution. As a benchmark we give in the first line the values for the fitted normal. It shows the poor performance of a normal fit.

In the next section we will compare empirical intraday distributions with members of the convolution semigroup of a fitted infinitely divisible distribution. Various

Name	$< 1\sigma$	$< 2\sigma$	$< 3\sigma$	$< 4\sigma$	$< 5\sigma$	$\geq 5\sigma$	
$N(0,1)$	0.683	0.954	0.997	1.000	1.0000	0.0000006	(0.001)
Bayer	0.756	0.946	0.988	0.996	0.9990	0.0010132	(2)
$H_{377.74,3.04,0.0017,0}$	0.750	0.947	0.989	0.998	0.9995	0.0004688	(1.04)
Daimler	0.743	0.945	0.987	0.998	1.0000	0.0000000	(0)
$H_{375.78,-7.49,0.0019,0}$	0.740	0.943	0.988	0.997	0.9995	0.0005140	(1.01)
Deutsche Bank	0.745	0.952	0.988	0.997	0.9990	0.0010209	(2)
$H_{519.99,-26.1,0.0016,0}$	0.747	0.948	0.990	0.998	0.9998	0.0003768	(0.74)
Siemens	0.747	0.948	0.990	0.999	0.9990	0.0010209	(2)
$H_{539.59,-25.56,0.0024,0}$	0.740	0.950	0.991	0.998	0.9997	0.0002627	(0.51)
Thyssen	0.750	0.943	0.987	0.997	0.9995	0.0005105	(1)
$H_{263.70,2.59,0.0016,0}$	0.751	0.943	0.987	0.997	0.9994	0.0006191	(1.21)

Table 2: Frequency distributions.

distances will be used. First we recall the Kolmogorov distance of distribution functions F_1, F_2

$$\mathrm{dist}_K(F_1, F_2) = \sup_{x \in \mathbb{R}} |F_1(x) - F_2(x)|.$$

Because of the 0-returns which occur for narrow time grids it makes sense to look also at the following modification of the classical Kolmogorov distance

$$\mathrm{dist}_{\widetilde{K}}(F_1, F_2, \varepsilon) = \sup_{|x| > \varepsilon} |F_1(x) - F_2(x)|.$$

If $f(0)$ is the relative frequency of 0-returns, the empirical distribution function has a jump of size $f(0)$ at the point $x = 0$. Consequently the Kolmogorov distance between the empirical and a fitted continuous distribution function is at least $\frac{1}{2}f(0)$. $\mathrm{dist}_{\widetilde{K}}$ eliminates this 0-return effect. Further we look at χ^2-distances, where we choose the classes such that each class has equal probability under F_1, and from the family of L_p distances

$$\mathrm{dist}_{L_p}(F_1, F_2) = \left(\int_{\mathbb{R}} |F_1(x) - F_2(x)|^p \, \mathrm{d}x \right)^{1/p}$$

we chose L_1. In the following Tables 3 and 4 the distance to the fitted normal is always given as a benchmark. dist_K and $\mathrm{dist}_{\widetilde{K}}$ produce the same values for the normal distribution.

In a refined intraday consideration we shall take the overnight returns into account. An overnight return is given by the difference of the log price at the beginning of day $j+1$ (opening price) minus the log price at the end of day j (closing price). Figure 5 shows the hyperbolic as well as the normal fit for overnight returns from Daimler data.

	dist_K for hyp. fit	$\text{dist}_{\widetilde{K}}$ for hyp. fit	$\text{dist}_K = \text{dist}_{\widetilde{K}}$ for normal fit
Bayer	0.03413	0.01795	0.05878
Daimler	0.03727	0.02155	0.05243
Deutsche Bank	0.03662	0.02155	0.05567
Siemens	0.03152	0.01966	0.04628
Thyssen	0.05844	0.02144	0.06218

Table 3: (Modified) Kolmogorov distance between the empirical distribution of log returns and the fitted hyperbolic resp. the fitted normal distribution.

	dist_{χ^2} for hyp. fit	dist_{χ^2} for normal fit	dist_{L_1} for hyp. fit	dist_{L_1} for normal fit
Bayer	72.11	193.66	0.00008	0.00049
Daimler	163.09	286.22	0.00020	0.00045
Deutsche Bank	386.42	548.01	0.00022	0.00037
Siemens	173.27	257.28	0.00018	0.00033
Thyssen	189.20	372.73	0.00015	0.00075

Table 4: χ^2- and L_1-distance between the empirical distribution of log returns and the fitted hyperbolic resp. the fitted normal distribution.

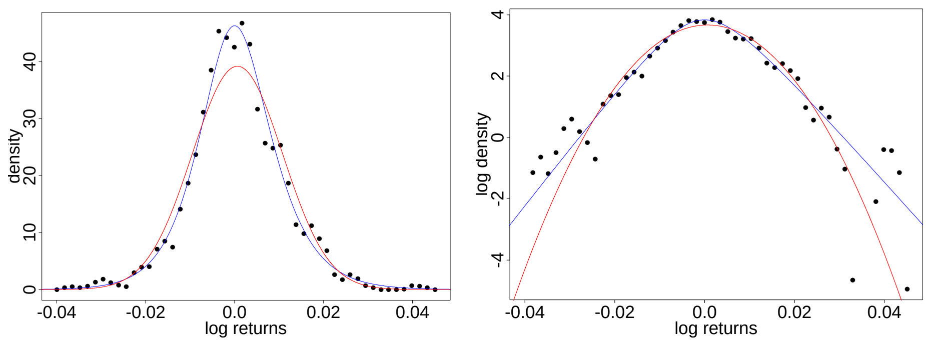

Figure 5: Empirical density of overnight returns (Daimler) vs. density of the fitted hyperbolic (blue) and the fitted normal distribution (red).

4 Downward and upward convolution

Now we will study the relationship between daily and intraday data. In the framework of Lévy processes and infinitely divisible distributions described above one could ask the question in the following form: Does the parameter t of the convolution semigroup correspond to (physical) time? There are two directions to consider. Starting with the calibration of the model on a daily basis one can look downwards to shorter time horizons. On the other side one can do the calibration on an intraday scale and then look upwards to larger time horizons. Of course once one has fixed a time grid for calibration, the model should behave well on all time horizons of interest in any direction. For simplicity we focus on daily and hourly data in this exposition.

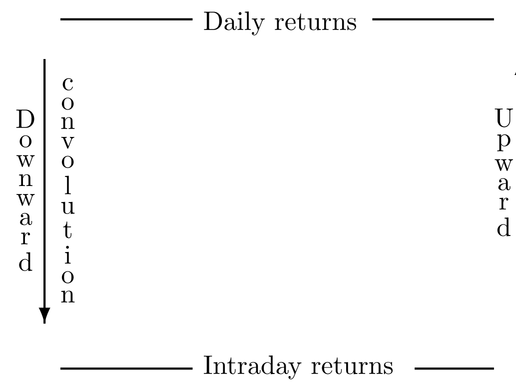

Figure 6: Relationship between the different time scales.

We start with the left side and fit daily Kassa returns by a hyperbolic distribution $\nu^K = H(\widehat{\alpha}, \widehat{\beta}, \widehat{\delta}, \widehat{\mu})$. Kassa returns, i.e. returns derived from prices recorded in the middle of the trading session, are appropriate since they are free of special opening and closing effects. ν^K generates a convolution semigroup $\left(\nu_t^K\right)_{t\geq 0}$ such that $\nu^K = \nu_1^K$. Figure 7 shows the distance between the elements of $(\nu_t^K)_{t\geq 0}$ and the empirical distribution of hourly returns of Bayer price quotes as a function of t. In the upper row on the left side we used the Kolmogorov distance, on the right side L_1-distances are plotted. At the bottom we see the corresponding χ^2-distances, where the left picture was obtained with classes of equal probability whereas the smoother right picture resulted from fixed classes for each element of the convolution semigroup.

The minimum is achieved for any of the distances around 0.3. This is an excellent value given the daily trading time of 3 hours. Table 5 shows that similar results for $t_{\min}$ are obtained for other equity values as well, where we take the modified Kolmogorov distance. Actually with one exception, namely Thyssen, $t_{\min}$ is the same for the classical and the modified Kolmogorov distance. The remaining three columns in the table give the modified Kolmogorov distances $\mathrm{dist}_{\widetilde{K}}$ of the empirical distribution of hourly returns to $\nu^K_{t_{\min}}$, to the fitted hyperbolic, and to the fitted normal distribution. As the table shows at least for Bayer, Daimler, and Thyssen the fit of the empirical distribution by the distribution derived from the convolution semigroup is better than the fit by the normal distribution. Let us complement this consideration by a graphical comparison. Figure 8 shows the empirical one-hour return distribution represented by the points and the element of the convolution semigroup derived from Kassa prices given by $t = 0.33$. On the right side the same is plotted on a log scale in order to exhibit the tail behavior. As an aside one can

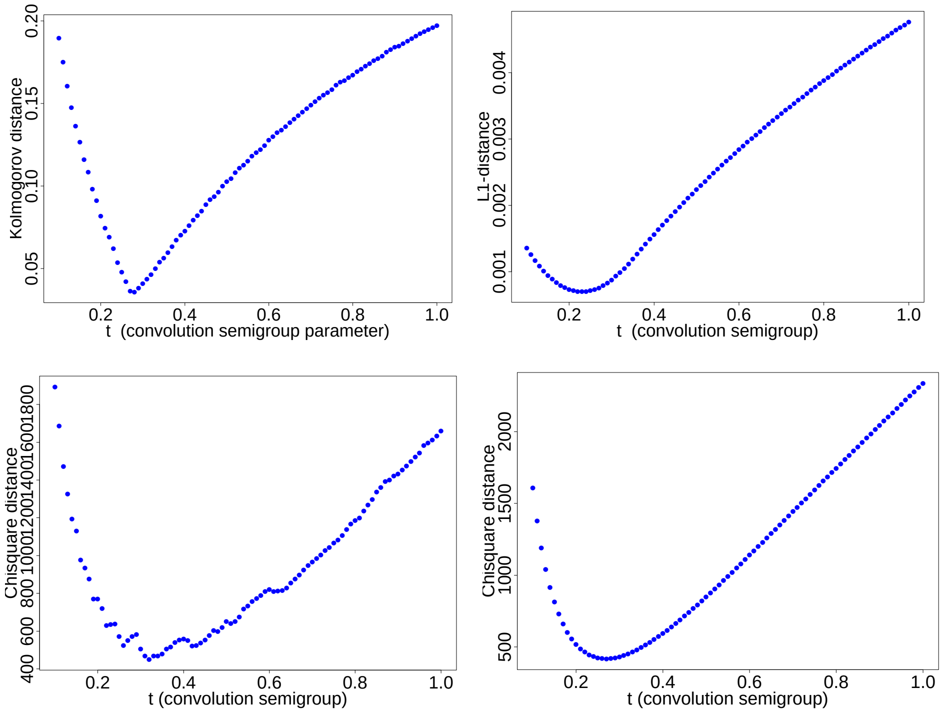

Figure 7: Distance between empirical density of hourly log returns (Bayer) and elements of the convolution semigroup $(\nu_t^K)_{0\leq t\leq 1}$.

Name	$t_{\min}$	$\nu^K_{t_{\min}}$	hyp. fit	norm. fit
Bayer	0.28	0.0358	0.0180	0.0588
Daimler	0.26	0.0524	0.0215	0.0524
Deutsche Bank	0.29	0.0850	0.0244	0.0557
Siemens	0.21	0.0600	0.0197	0.0463
Thyssen	0.30	0.0414	0.0214	0.0622

Table 5: Modified Kolmogorov distances.

see from this log density that the elements of the convolution semigroup ν_t^K $(t \neq 1)$ are no longer hyperbolic. Otherwise one would see a hyperbola. The tails of $\nu_{0.33}^K$ are heavier.

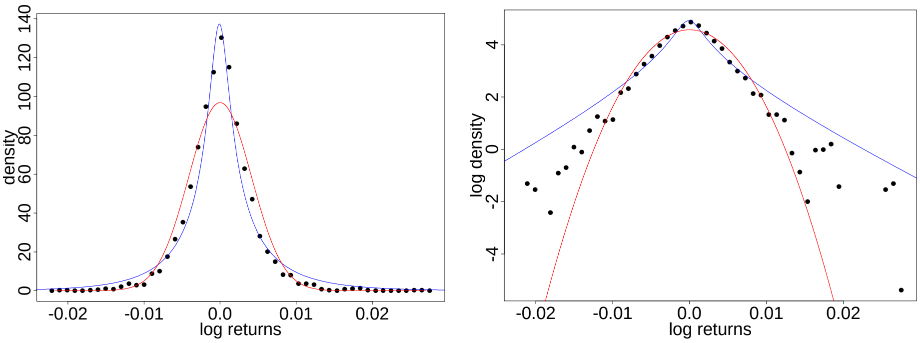

Figure 8: Empirical density of one-hour returns (Bayer) vs. density of $\nu_{0.33}^K$ (blue) and fitted normal distribution (red).

So far our focus was on downward convolution from a daily to a 60-minutes time grid. The same analysis can be done for other time grids as well. In the following table we show the results for a number of m-minutes returns, $m = 10, 15, \dots, 90$ for Daimler data.

m	10	15	20	30	45	60	90
$\mathrm{t_{min}}$	0.12	0.14	0.16	0.19	0.22	0.26	0.32
$\mathrm{dist}_{\widetilde{K}}(F_{m;\mathrm{intra}}, \nu^K_{\mathrm{t_{min}}})$	0.075	0.077	0.071	0.066	0.062	0.052	0.049

Table 6: Different time grids: $\mathrm{t_{min}}$ and $\mathrm{dist}_{\widetilde{K}}(F_{m;\mathrm{intra}}, \nu^K_{\mathrm{t_{min}}})$ (Daimler).

Again the empirical distribution of m-minutes intraday returns, denoted by $F_{m;\mathrm{intra}}$, is compared to $\left(\nu_t^K\right)_{0 \leq t \leq 1}$. We derive $\mathrm{t_{min}}$ and the modified Kolmogorov distance to $\nu^K_{\mathrm{t_{min}}}$. The results for other equities are rather similar and are fairly independent of the distance chosen. It is not surprising that as m becomes smaller it is more difficult to get a good fit from the convolution semigroup.

Before we turn to upward convolution we have a closer look at overnight returns. The optimal element from $(\nu_t^K)_{0 \leq t \leq 1}$ to fit hourly returns was always obtained for a t somewhat below 0.33. This is because we ignored overnight returns. Therefore it is natural to subtract the overnight returns, which we shall denote by Y, from the daily Kassa returns. We show the effect in Figure 9. The left side are intraday Bayer quotes for the period January 2 to January 7, 1992, the right side shows the quotes with the overnight jumps removed.

Again we fit a hyperbolic distribution to the modified data set $K - Y$ and get the corresponding convolution semigroup $\left(\nu_t^{K-Y}\right)_{t \geq 0}$ whose elements are compared with the empirical distribution of hourly returns. As Table 7 shows, the minimal distance is achieved now for values of t clearly above 0.33.

In Figure 10 we show the excellent fit which $\nu_{0.42}^{K-Y}$ from the Bayer convolution semigroup provides for the empirical hourly returns.

Now we turn to the right side of the scheme in Figure 6, the upward convolution. Calibration is done on an hourly basis where to be concise we immediately take

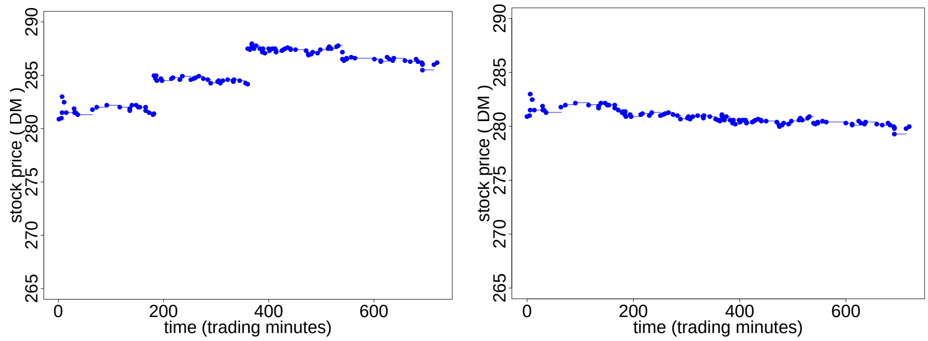

Figure 9: Bayer intraday quotes for the period January 2 to January 7, 1992, the right side shows the quotes with the overnight jumps removed.

Name	$\mathrm{t_{min}}$ for $\mathrm{dist}_{\widetilde{\mathrm{K}}}$	$\mathrm{t_{min}}$ for dist_{χ^2}
Bayer	0.42	0.47
Daimler	0.51	0.59
Deutsche Bank	0.45	0.39
Siemens	0.40	0.37
Thyssen	0.45	0.56

Table 7: $\mathrm{t_{min}}$ for $(\nu_t^{K-Y})_{t\geq 0}$.

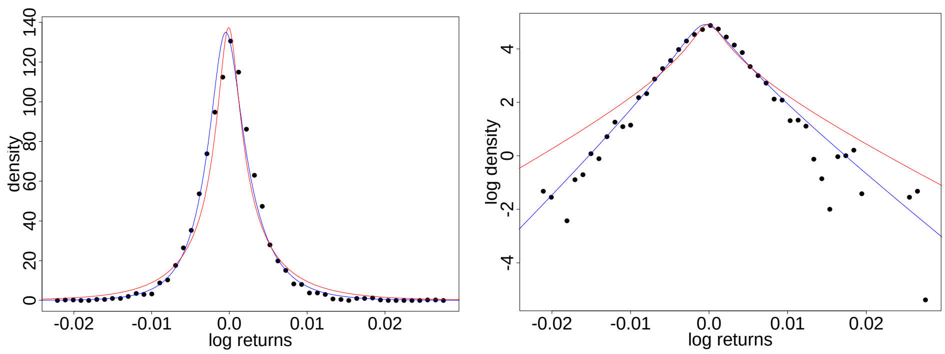

Figure 10: Empirical density of hourly returns (Bayer) vs. $(\nu_{0.42}^{K-Y})$ (blue) and $(\nu_{0.33}^{K})$ (red).

overnight returns into account. Since we compare the result of upward convolution to daily Kassa returns we define

$$\begin{aligned} F_{\text{conv60}\uparrow} &:= F_{\text{hyp.60min}} * F_{\text{hyp.60min}} * F_{\text{hyp.60min}} * F_{\text{hyp.nacht}} \\ &= (F_{\text{hyp.60min}})^{*3} * F_{\text{hyp.nacht}} \end{aligned} \tag{4.1}$$

where $F_{\text{hyp.60min}}$ denotes the hyperbolic distribution fitted to one hour returns and $F_{\text{hyp.nacht}}$ denotes the hyperbolic distribution fitted to overnight returns. In Figure 11 we show for Daimler data the empirical density as points. The density of the upward convolution defined in (4.1) is given as a blue line. The hyperbolic fit of Kassa returns is given as a red line.

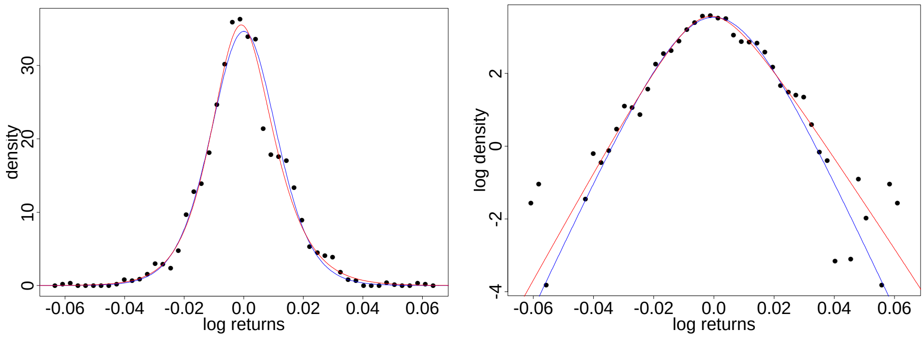

Figure 11: Empirical density of daily Kassa returns (Daimler) vs. $F_{\text{conv60}\uparrow}$ (blue) and fitted hyperbolic distribution (red).

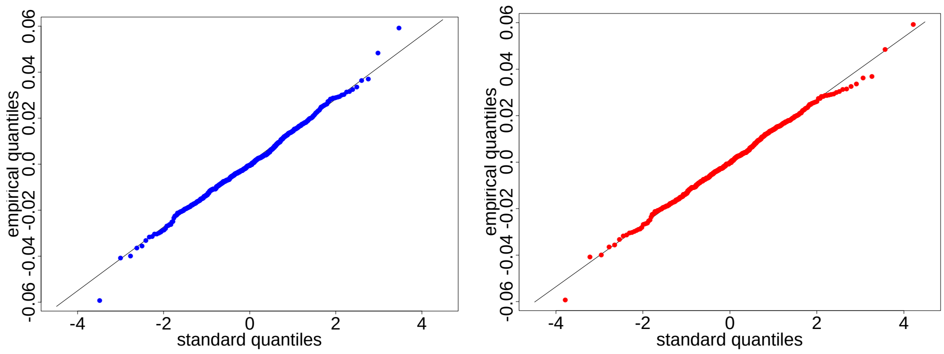

Figure 12: QQ-plot of daily Kassa returns (Daimler) vs. $F_{\text{conv60}\uparrow}$ (blue) and fitted hyperbolic distribution (red).

Figure 12 shows the QQ-plots empirical Kassa versus $F_{\text{conv60}\uparrow}$ on the left side and empirical Kassa versus fitted hyperbolic on the right side. Essentially the same quality of the fit is achieved if one starts with fitting thirty-minutes returns and compares

$$F_{\text{conv30}\uparrow} := (F_{\text{hyp.30min}})^{*6} * F_{\text{hyp.nacht}}$$

with Kassa returns. Figure 13 shows the QQ-plot for Daimler data.

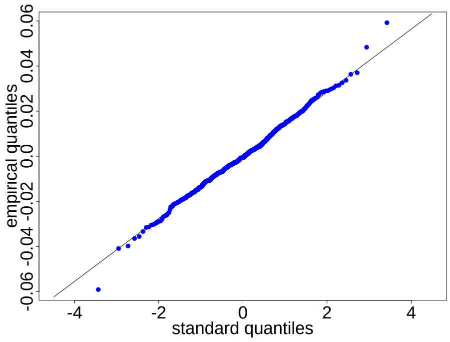

Figure 13: QQ-plot of daily Kassa returns (Daimler) vs. $F_{\mathrm{conv}30\uparrow}$.

	dist_K for hyp. fit	dist_K for $F_{\mathrm{conv}60\uparrow}$	dist_{χ^2} for hyp. fit	dist_{χ^2} for $F_{\mathrm{conv}60\uparrow}$
Bayer	0.01411	0.04349	29.35	36.30
Daimler	0.02389	0.03257	35.15	36.94
Deutsche Bank	0.02020	0.03318	24.45	26.28
Siemens	0.02461	0.02813	21.09	20.91
Thyssen	0.02116	0.05424	30.09	35.99

	dist_{L_1} for hyp. fit	dist_{L_1} for $F_{\mathrm{conv}60\uparrow}$	dist_{L_2} for hyp. fit	dist_{L_2} for $F_{\mathrm{conv}60\uparrow}$
Bayer	0.000279	0.000646	0.001185	0.003442
Daimler	0.000429	0.000644	0.001917	0.002896
Deutsche Bank	0.000334	0.000497	0.001419	0.002117
Siemens	0.000345	0.000335	0.001566	0.001598
Thyssen	0.000590	0.000964	0.002168	0.004252

Table 8: Distribution distances between empirical distribution of daily Kassa returns vs. hyperbolic fit and $F_{\mathrm{conv}60\uparrow}$.

For a more quantitative comparison we list in Table 8 Kolmogorov-, χ^2-, L_1-, and L_2-distances. The classes for the χ^2-distance were built according to the rule of Moore.

We also tested the hypothesis H_0 : "daily Kassa returns are $F_{\text{conv60}\uparrow}$-distributed" against H_1 : "daily Kassa returns are not $F_{\text{conv60}\uparrow}$-distributed". Table 9 shows the results of three different χ^2-tests on a level $\alpha = 0.01$. For $\widehat{\chi}_1^2$ the number of classes ($k = 23$) was determined according to the rule of Moore, $\widehat{\chi}_2^2$ ($k = 43$) is based on the rule of Mann and Wald. For $\widehat{\chi}_3^2$ we started with $k = 80$ classes of equal probability and then merged classes such that for each class the number of expected observations is at least five. For the choice of the degrees of freedom of the χ^2-distribution see e.g. Rice (1988). Only $\widehat{\chi}_3^2$ for Daimler exceeds the critical value $\chi^2_{k-1;0.99}$. In all other cases the test accepts.

Name	$\widehat{\chi}_1^2$	$\chi^2_{k-1;0.99}$	$\widehat{\chi}_2^2$	$\chi^2_{k-1;0.99}$	$\widehat{\chi}_3^2$	$\chi^2_{k-1;0.99}$
Bayer	36.30	45.64	51.13	66.21	89.95	111.14
Daimler	36.94	45.64	46.28	66.21	119.61	111.14
Deutsche Bank	26.28	45.64	57.58	66.21	108.92	111.14
Siemens	20.91	45.64	37.75	66.21	81.76	111.14
Thyssen	35.99	45.64	60.7	66.21	108.29	111.14

Table 9: χ^2-tests for H_0 : "daily Kassa returns are $F_{\text{conv60}\uparrow}$-distributed" against H_1 : "daily Kassa returns are not $F_{\text{conv60}\uparrow}$-distributed"

An analysis similar to the one exposed here for equity price data, was made for bond prices as well. The underlying model is the Lévy term structure model developed in Eberlein and Raible (1999) and Eberlein and Özkan (2003). The model was fitted to daily bond prices and we compared the upward convoluted distribution on a 3-day, 5-day (weekly) and 10-day (2 weeks) horizon to the corresponding empirical return distributions. Again the Lévy model turned out to be highly consistent.

Acknowledgement: We want to thank Deutsche Börse AG for providing high-frequency stock price data for a number of values traded at the Frankfurt stock exchange. We also thank two referees for corrections and suggestions which helped to improve the paper.

References

Barndorff-Nielsen, O. E. (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society London A 353*, 401 – 419.

Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type. *Finance and Stochastics 2*, 41 – 68.

Chan, T. (1999). Pricing contingent claims on stocks driven by Lévy processes. *The Annals of Applied Probability 9*, 504 – 528.

Eberlein, E., J. Kallsen, and J. Kristen (2003). Risk management based on stochastic volatility. To appear in *Journal of Risk*.

Eberlein, E. and U. Keller (1995). Hyperbolic distributions in finance. *Bernoulli 1*, 281 – 299.

Eberlein, E. and F. Özkan (2003). The defaultable Lévy term structure: Ratings and restructuring. To appear in *Mathematical Finance*.

Eberlein, E. and K. Prause (2002). The generalized hyperbolic model: financial derivatives and risk measures. In *Mathematical Finance – Bachelier Congress 2000*, pp. 245 – 267.

Eberlein, E. and S. Raible (1999). Term structure models driven by general Lévy processes. *Mathematical Finance 9*, 31 – 53.

Eberlein, E. and E. A. v. Hammerstein (2002). Generalized hyperbolic and inverse Gaussian distributions: Limiting cases and approximation of processes. Preprint University of Freiburg.

Jacod, J. and A. N. Shiryaev (1987). *Limit Theorems for Stochastic Processes*. Springer.

Madan, D. B. and F. Milne (1991). Option pricing with V.G. martingale components. *Mathematical Finance 1*, 39 – 56.

Raible, S. (2000). *Lévy processes in finance: theory, numerics, and empirical facts*. Ph. D. thesis, University of Freiburg.

Rice, J. A. (1988). *Mathematical Statistics and Data Analysis*. Wadsworth & Brooks/Cole Advanced Books & Software.

Sato, K.-I. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press.

Skorohod, A. V. (1991). *Random Processes with Independent Increments*. Kluwer Academic Publishers.