

QUANTIFICATION OF RISK IN NORWEGIAN STOCKS VIA THE NORMAL INVERSE GAUSSIAN DISTRIBUTION

ERIK BØLVIKEN AND FRED ESPEN BENTH

ABSTRACT. The quantification of risk in Norwegian stocks via the normal inverse Gaussian distribution is studied. On seven stocks quoted on the Norwegian Stock Exchange and one quoted in New York we estimate Value-at-Risk on different time horizons for the fitted Gaussian distribution, the normal inverse Gaussian distribution and a non-parametric model. Results are compared and discussed.

1. INTRODUCTION

The family of normal inverse Gaussian distributions, constructed more than two decades ago by Barndorff-Nielsen [2], has the remarkable property of being closed under convolutions, yet at the same time able to portray stochastic phenomena that have heavy tails or are strongly skewed. Unlike the gamma family it is not confined to the positive half axis. The model is an obvious candidate for financial data, where experience dictates that the Gaussian family often underestimates random variation, even after passing to logarithms of, say prices of stocks and other securities. With the normal inverse Gaussian distributions the financial analyst has at its disposal a model that can flexibly be adapted to many different shapes while the distribution of sums of independent random variables are still trivial to compute.

Applications in finance have been reported in several papers, for example Barndorff-Nielsen [3], Rydberg [14] and Prause [11]. The purpose of this article is to investigate this model as a tool to evaluate the uncertainty in future prices of stock listed on the Norwegian Stock Exchange in Oslo. For comparison we shall also consider a Norwegian security quoted in New York. Chiefly among the questions to be addressed is how well the normal inverse Gaussian family fits share prices at one of the smaller financial communities of the world where trading is relatively low and where the flow of information may differ in nature from what it is at larger financial centres? How much bias is introduced by postulating that the day to day fluctuations follow a normal inverse Gaussian distribution, which can not possibly be correct? How much data are needed to estimate the parameters of normal inverse Gaussian models sufficiently accurate for Value-at-risk evaluations?

Date: February 16, 2000.

Our framework are perfect markets where the changes from one day to the next are independent of what has happened before. This means that we are imposing Lévy processes or, in discrete time, random walk models. A brief discussion of these models is given in Section 2.1 below. Note the usefulness of the convolution property in this context. If Value-at-Risk are sought over longer time horizons, we can simply plug in the appropriate member of the normal inverse Gaussian family, with parameters depending on the time interval in question. Other ways of evaluating Value-at-Risk will have to perform some sort of numerical integration to achieve the same thing. It is outside our brief to compare the normal inverse Gaussian approach to other proposals in the literature, for example the ideas in Embrechts, Klüppelberg and Mikosch [9], who employ Extreme Value Theory and generalised Pareto distributions. In the next section the normal inverse Gaussian family of models and its statistical inference will be introduced in a tutorial manner. Value-at-Risk evaluation will be considered in Section 3, and eight different Norwegian securities studied in Section 4.

2. MATHEMATICAL AND STATISTICAL MODELLING

2.1. Lévy processes. A Lévy process L_t is a time-continuous stochastic process with independent and stationary increments; for every $t > s \geq 0$, $L_t - L_s$ is independent of L_s and its distribution depends only on the time increment $t - s$ and not on t or s . Classical examples are Brownian motion and the Poisson process. Except for Brownian motion, the paths of a Lévy process have jumps at random time points, and that is precisely what goes on in a market of securities if we take the view that the value is the price agreed last. But the Lévy model also imposes serious restrictions. The sizes of the jumps, whether positive or negative for example, have no relation to the price level reached, and the continuous time viewpoint forces the distribution of the increments $L_t - L_s$ to belong to the infinitely divisible ones; e.g the distribution must for every n be the n -fold convolution of some other distribution. There are not too many models with this property around, but the normal inverse Gaussian family is one of them.

Arguably there is no lower limit as to how close in time two different tradings can be completed, but in practice it is common to pass to a discrete viewpoint and track the stock price at some sequence of equidistant points in time. If a very fine resolution is used (say seconds!) the distribution of the increment $L_t - L_s$ will have a huge atom at the origin ('no trading'), but when we deal in terms of days and liquid shares, a purely continuous distribution, without atoms, may be in order. Experience also dictates that it is reasonable to relate the stock price S_t to the Lévy process L_t through

$$(2.1) \quad S_t = S_0 \exp(L_t),$$

known as a *Geometric Lévy* process. The ratio of prices at t and $t - \Delta$ now becomes

$$(2.2) \quad \ln(S_t/S_{t-\Delta}) = L_t - L_{t-\Delta},$$

and we pass in a discrete time to an ordinary random walk based on independent increments $X_t = L_t - L_{t-\Delta}$, the distribution of which being our modelling tool. In this paper the distribution of X_t will be assumed to belong to a specific parametric form.

2.2. The normal inverse Gaussian distribution. A random variable X follows a normal, inverse Gaussian distribution with parameter vector $(\mu, \delta, \alpha, \beta)$, in symbolic notation $X \sim \text{NIG}(\mu, \delta, \alpha, \beta)$, if its probability density function is

$$(2.3) \quad \text{nig}(x; \mu, \delta, \alpha, \beta) = \delta^{-1} \text{nig}\left(\frac{x - \mu}{\delta}; 0, 1, \alpha, \beta\right)$$

where

$$(2.4) \quad \text{nig}(z; 0, 1, \alpha, \beta) = \frac{\alpha}{\pi} \exp(\sqrt{\alpha^2 - \beta^2} + \beta z) \frac{K_1(\alpha\sqrt{1+z^2})}{\sqrt{1+z^2}}.$$

Here K_1 is a modified Bessel function of the third kind with index 1; see Section 2.3 below. Note that μ and δ are ordinary parameters of location and scale whereas α and β determines the *shape* of the density; i.e. if $X \sim \text{NIG}(\mu, \delta, \alpha, \beta)$, then $Z = (X - \mu)/\delta$ is the *standardized* version $\text{NIG}(0, 1, \alpha, \beta)$. We have chosen to parametrize the family as suggested in Barndorff-Nielsen [3], but different from many other papers, notably Rydberg [14], where α and β have a slightly different meaning. It is in our view more transparent to let the shape of the density be determined by exactly two of the four parameters. The conditions for a viable density is $\delta > 0$, $\alpha > 0$ and $|\beta|/\alpha < 1$.

The mean, variance, skewness and kurtosis of X are

$$\begin{aligned} E[X] &= \mu + \delta \frac{\beta/\alpha}{(1 - (\beta/\alpha)^2)^{1/2}} \\ \text{Var}[X] &= \delta^2 \alpha^{-1} \frac{1}{(1 - (\beta/\alpha)^2)^{3/2}} \\ \text{Skew}[X] &= 3\alpha^{-1/4} \frac{\beta/\alpha}{(1 - (\beta/\alpha)^2)^{1/4}} \\ \text{Kurt}[X] &= 3\alpha^{-1/2} \frac{1 + 4(\beta/\alpha)^2}{(1 - (\beta/\alpha)^2)^{1/2}} \end{aligned}$$

Here, $\text{Skew}[X] = E[(X - E[X])^3]/(\text{Var}[X])^{3/2}$ and $\text{Kurt}[X] = E[(X - E[X])^4]/(\text{Var}[X])^2 - 3$. Note that

$$\frac{\text{Skew}[X]^2}{\text{Kurt}[X]} = \frac{3(\beta/\alpha)^2}{1 + 4(\beta/\alpha)^2},$$

so that

$$\text{Kurt}[X] > 0 \text{ and } |\text{Skew}[X]| \leq \frac{3}{5} \text{Kurt}[X],$$

(since $|\beta|/\alpha < 1$). There is thus a bound on the skewness relative the kurtosis, and the tails of the distribution is always heavier than those of the normal (which has zero kurtosis). In fact, (see e.g. [3]),

$$\text{nig}(x; \alpha, \beta, \mu, \delta) \sim k|x|^{-3/2} \exp(-(\alpha/\delta)|x| + (\beta/\delta)x), \text{ when } |x| \rightarrow \infty,$$

where k is a positive constant. Heavy tails is hardly any practical restriction in empirical finance. The distribution converges to the normal as $\alpha \rightarrow \infty$ (if $\delta \propto \alpha^{1/2}$) whereas the Cauchy distribution appears as the opposite limit as $\alpha \rightarrow 0$. The parameter β determines the degree of skewness. For symmetrical densities $\beta = 0$.

An important property of the the family of normal inverse Gaussian distributions is its behaviour under convolutions. Suppose X_1 and X_2 are independent and distributed as $\text{NIG}(\mu_1, \delta_1, \alpha_1, \beta_1)$ and $\text{NIG}(\mu_2, \delta_2, \alpha_2, \beta_2)$ respectively. Then, $X_1 + X_2 \sim \text{NIG}(\mu_1 + \mu_2, \delta_1 + \delta_2, \alpha_1 + \alpha_2, \beta_1 + \beta_2)$, provided $\alpha_1/\delta_1 = \alpha_2/\delta_2$. In particular, any sum of independent and identically NIG-distributed random variables is also NIG-distributed. By appealing to the central limit theorem the result underpins the earlier claim that the normal distribution will eventually appear, as α grows. If logreturns follow a random walk, as suggested in Section 2.1, see (2.1) in particular, one would expect logreturn data to behave more like normal ones over longer time spans. For example, Eberlein and Keller [6] conclude that the normal distribution is appropriate for monthly logreturn data, but reject this distribution for daily or weekly logreturn data¹.

The normal inverse Gaussian distribution can be generalised with a fifth parameter to the so-called generalized inverse Gaussian distributions; see e.g. Rydberg [14]. Eberlein and Keller [6] used a subfamily called the hyperbolic distributions to study logreturn data from the German stock market (see also [7] and [12]).

2.3. Fitting normal inverse Gaussian models. It is easy to see from the preceding discussion that the expressions for $E[X]$, $\text{Var}[X]$, $\text{Skew}[X]$ and $\text{Kurt}[X]$ can be inverted to yield a unique set of parameters $(\mu, \delta, \alpha, \beta)$. For example, the ratio β/α is determined from $(\text{Skew}[X])^2/\text{Kurt}[X]$, and the others then follow readily. This means that the so-called method of moments is available to estimate the parameters from a given sample x_1, \dots, x_n of normal inverse Gaussian distributed random variables. The mean, variance, skewness and kurtosis are then replaced by their sample versions and the four equations are solved for α, β, μ and δ . Such equations are unwieldy in many other situations in statistics, but here they are straightforward, and the whole operation is numerically so simple that the Bessel function does not have to be evaluated at all.

¹Their study was performed on German stocks. See the references in [6] for empirical investigations of US stocks

There is, however, reason to be suspicious of the quality of such estimates based on higher order sample moments, and there is, of course, no guarantee that these sample moments satisfy the restrictions laid down by the normal inverse Gaussian family at all, so that the moment equations have actually a solution. Likelihood estimation seems a better, albeit more complicated alternative. This method was suggested in Blæsild and Sørensen [5] and Rydberg [14]. The former used the method of steepest descent to maximize the likelihood surface whereas the latter found a zero of the differentiated likelihood functions. The numerical literature (for example Gill, Murray and Wright [10]) usually recommends other approaches. Prominent are the two classes of optimization methods known as the variable metrics and the conjugated gradients. We have tested both on the data in the next section. Both worked well when properly implemented, but the variable metric method (BFGS version) converged faster and is our recommendation.

Either of these approaches require the evaluation of the log-likelihood function

$$(2.5) \quad \ell(\alpha, \beta, \mu, \delta) = \sum_{i=1}^n \log(\text{nig}(x_i; \mu, \delta, \alpha, \beta))$$

and its partial derivatives with respect to the parameters. The evaluation of the Bessel function $K_1(z)$ and its derivative is then required. There exist countless representations of these functions; see Section 9.6 of Abramowitz and Stegun [1], for example,

$$(2.6) \quad K_1(z) = z \int_1^\infty e^{-zt} \sqrt{t^2 - 1} dt$$

from which it can be proved that

$$(2.7) \quad K_1'(z) = -z^{-1}K_1(z) - z^{-2}K_0(z)$$

where $K_0(z) = \int_1^\infty e^{-zt}(t^2 - 1)^{-1/2} dt$. Subroutines for the evaluation of K_0 and K_1 are available in the Numerical Recipe library, see Press, Teukolsky, Vetterling and Flannery [13].

3. VALUE-AT-RISK

We define the concept of Value-at-Risk according to JP Morgans RiskMetrics: The logreturn from the stock S_t at time t is $L_t = \ln(S_t/S_0)$. If p is a specified level risk tolerance, the Value-at-Risk for S_t at time t is

$$(3.1) \quad \text{VaR}_p(S_t) = E[L_t] - q_p(t)$$

where $q_p(t)$ is the p -quantile of L_t . In the standard stock price model, $L_t = \gamma t + \sigma B_t$ for a Brownian motion B_t with drift γ and volatility σ . In this case we recover the well-known formula for Value-at-Risk

$$(3.2) \quad \text{VaR}_p(S_t) = \epsilon_p \sigma \sqrt{t}$$

where ϵ_p is the p -quantile of a standard normal distribution. It is of interest to compare the Value-at-Risk estimate in the standard model with that of a normal inverse Gaussian model. We expect the two estimates to converge for large time horizons. To explain this, let L_1 be normal inverse Gaussian distributed with parameters α, β, μ and δ . Then by the convolution property the logreturn at time t is distributed as $L_t \sim \text{nig}(x; \alpha t, \beta t, \mu t, \delta t)$, and thus

$$\begin{aligned} E[L_t] &= \left(\mu + \frac{\delta\beta}{\sqrt{\alpha^2 - \beta^2}} \right) \cdot t \\ \text{Var}[L_t] &= \frac{\delta^2 \alpha^2}{(\sqrt{\alpha^2 - \beta^2})^3} \cdot t \end{aligned}$$

When t increases, we know that L_t tends to a normal distribution. In fact for large t , we approximately have

$$L_t \sim \mathcal{N}\left(\left(\mu + \frac{\delta\beta}{\sqrt{\alpha^2 - \beta^2}}\right) \cdot t, \frac{\delta^2 \alpha^2}{(\sqrt{\alpha^2 - \beta^2})^3} \cdot t\right)$$

Hence, if we are able to fit perfectly both a normal inverse Gaussian and a Gaussian model to the same set of data, the corresponding Value-at-Risk estimates will be nearly identical for large time horizon. However, on shorter time scales there will be a significant difference between the two Value-at-Risk estimates.

4. EMPIRICAL STUDY

4.1. Summary description and fit. We have used eight Norwegian stocks as basis for an empirical evaluation of the normal inverse Gaussian model. Seven of them (Agresso, Bergesen, Den norske Bank, Kreditkassen, Merkantildata, Norsk Hydro and Petroleum Geo-Services) are quoted on the Oslo Stock Exchange while the eight (also Norsk Hydro, but now denoted Norsk Hydro (NYSE)) was taken from the notations in New York. The returns from the former group of seven papers were end-of-the-day prices from October 16, 1997 until November 17, 1999, that is 506 different values. The data from New York was longer, lasting from January 2, 1990 until December 31, 1998, a total of 2274 notations, all at closing time. When dealing with increments, all series become one unit shorter. All analyses below are in terms of log-returns and assume the stock prices to follow a random walk.

The first part of the analysis is a summary description of the data and an evaluation of the fit. The estimates of the parameters are shown in Table 1 with estimates of their standard error obtained by the bootstrap technique explained in 4.2 below. We have in Figure 1 converted the estimates of the shape parameters α and β to an alternative, more pictorial form sometimes used. Introduce

$$(4.1) \quad \xi = (1 + \sqrt{\alpha^2 - \beta^2})^{-1/2}, \quad \chi = \frac{\beta}{\alpha} \xi.$$

Here (χ, ξ) $0 \leq |\chi| < \xi < 1$ (since $|\beta| < \alpha$), and the line $\chi = 0$ divides the triangle in two parts according to the orientations of skewness, distributions skewed to the right (left) being on the right (left) hand side. The vertical middle line $\chi = 0$ (where $\beta = 0$) represents symmetric distributions, the normal being approached as $\xi \rightarrow 0$ and the Cauchy as $\xi \rightarrow 1$. The general picture is that of a phenomenon with considerably higher variability than the normal (ξ being a good deal greater than 0) whereas a possible skewness in the original increments evidently has been removed by passing to logarithms. The variability in ξ among the various papers reflect an underlying volatility that seems to differ between some of the shares.

This comes out in Figure 2, as well, where the estimated normal inverse Gaussian densities have been plotted against non-parametric estimates (obtained from the kernel method with a Gaussian kernel; see Silverman [15]). Note that the scale of the vertical axis has been made logarithmic, so that the normal density would come out as a parabola. The fit of the fitted normal inverse Gaussian distributions is excellent everywhere. The discrepancies at the extreme tails can be explained by larger random error of the non-parametric estimate in these regions.

4.2. Value-at-Risk. Using normal inverse Gaussian models to evaluate Value-at-Risk raises the objection that the estimated model is a compromise fit over the whole range of variation whereas value at risk deals with the left tail. However, with time horizons longer than one time unit, for example 5 or 10 days, as below, one does need the whole distribution to compute the five-fold or ten-fold convolution. The question is how much bias the normal inverse Gaussian assumption will bring and how large this effect is compared to the random variation. This, in turn, depends on how long the historical records are; i.e. how far back experience is deemed to be of relevance for the uncertainty of the future. The answer is, of course, also influenced by the percentile in question.

Consider Value-at-Risk at 1% and compare evaluations based on the normal family, the normal inverse Gaussian family and a purely non-parametric approach. The latter uses, for each stock, the empirical distribution function \hat{F} of the observed increments of logreturns as the estimate of the underlying distribution function F . The precise definition of \hat{F} is as the distribution function assigning probability mass n^{-1} to each observed difference x_j . To compute Value-at-risk t days ahead under this distribution, t increments were drawn with equal probabilities and *with replacement* from the set $\{x_1, \dots, x_n\}$. Their sum is then a realisation of the random walk over t days under \hat{F} . When this operation is repeated a large number N times, the 1% smallest is an approximation to the lower 1% percentile. The results reported are based on $N = 100000$ times repetitions, which is large enough for the impact of the Monte Carlo error to be negligible. The method was used even for $t = 1$ as a convenient device to circumvent that the exact 1% percetile of \hat{F} is not precisely defined.

The uncertainty of all the Value-at-Risk estimates was assessed by non-parametric bootstrapping; see Effron and Tibshriani [8]. Artificial data are then sampled from the estimate \hat{F} of the true distribution function F and the computations leading from the real data x_1, \dots, x_n to the estimate of the Value-at-Risk is copied *exactly* on the Monte Carlo data. When the latter operation is repeated, we can estimate the standard error of the Monte Carlo estimates and obtain a picture of the uncertainty of the original estimate. The method has the advantage that no assumption on the underlying distribution is imposed. With the Model Carlo estimate the approach implies a nested sampling scheme, but sampling from \hat{F} is so fast that the computer time is still minor. The number of Monte Carlo samples was 50.

The Value-at-Risk estimates are shown in Table 2. Note that the normal ones are consistently lower than the two others, consistent with the tails in Figure 1 being heavier than the normal. There is no consistent pattern of difference between the two other estimates, and this supports the normal inverse Gaussian model as suitable for Value-at-Risk evaluations. Note, however, the much higher variability of the non-parametric estimates. We have not investigated the reason for the uncertainty of the latter varying so little with the time horizon t , but it seems that the underlying exact function must grow very slowly.

A final aspect is how long historical records must be for the estimation of the parameters of a normal inverse Gaussian model to be sufficiently accurate for the Value-at-Risk estimate to be a reliable guide of future uncertainty. To obtain a partial answer we divided the available data into periods about a half year each (corresponding to around 125 logreturns increments). That gave four different set of notations for each of the seven stocks from the Oslo Stock Exchange and twenty for the one from New York. The corresponding Value-at-Risk estimates are plotted jointly in Figure 3 for a period of up to a month. The dotted curves are based on all available data; e.g two years for the Norwegian notations.

The results were not quite as expected. In all of the seven Norwegian cases the Value-at-Risk assessments based on half year data are consistently lower than the curves based on all data. There is no bug! The likelihood estimates of the parameters of the normal inverse Gaussian distributions were carefully checked to represent (local) maxima on the log likelihood surface with zero derivatives (in some cases the iteration converged to an infinite α corresponding to a Gaussian model). We also tried to find other local maxima by starting the iteration elsewhere, but without success. In other words, this rather peculiar phenomenon is real and is caused by huge upward bias in the estimates of α for half-year data, which undervalues the uncertainty of the stock. Note the completely different behaviour of the notations in New York. It is tempting to speculate that there could be some systematic differences between the two sites, but the results need to be backed up by Monte Carlo studies based on artificial data.

QUANTIFICATION OF RISK IN NORWEGIAN STOCKS

Stock	μ	δ	α	β/α
Agresso	-0.0028 (0.003)	0.033 (0.010)	0.64 (0.6)	0.047 (0.09)
Bergesen	-0.0067 (0.003)	0.025 (0.004)	1.25 (0.6)	0.210 (0.09)
Den norske Bank	0.0005 (0.001)	0.019 (0.002)	0.55 (0.2)	-0.028 (0.09)
Kreditkassen	0.0001 (0.001)	0.020 (0.002)	0.49 (0.2)	0.028 (0.09)
Merkantildata	-0.0031 (0.003)	0.034 (0.005)	0.85 (0.3)	0.018 (0.09)
Norsk Hydro	-0.0039 (0.009)	0.034 (0.025)	2.81 (9.0)	0.096 (0.12)
Norsk Hydro (NYSE)	-0.0006 (0.003)	0.015 (0.001)	0.86 (0.1)	0.047 (0.05)
Petroleum Geo services	-0.0048 (0.004)	0.041 (0.007)	1.07 (0.5)	0.083 (0.08)

TABLE 1. Estimated parameters in the normal inverse Gaussian distribution for eight Norwegian stock; estimated standard error in parenthesis

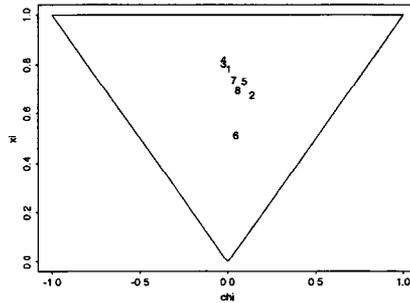


FIGURE 1. The estimated normal inverse Gaussian distributions plotted in the shape triangle. The numbers corresponds to the ordering of the stock in Table 1.

5. CONCLUSION

The normal inverse Gaussian distribution is a very flexible family of distributions enjoying the convolution property. With only four parameters it captures stylized facts like heavy tails and skewness observed in the marginals of financial time series data. By using sophisticated numerical optimization techniques, likelihood estimation of parameters in the normal inverse Gaussian distribution can be calculated efficiently and reliable, making it admissible for financial applications. We have demonstrated for seven stocks quoted on the Norwegian Stock Exchange and one quoted in New York that the normal inverse Gaussian distribution fits logreturn data nearly perfectly both in the tails and in the center, clearly outperforming the normal distribution.

Value-at-Risk with 1% risk tolerance was estimated for all eight stocks using the normal inverse Gaussian, normal and a non-parametric model. The comparison of the

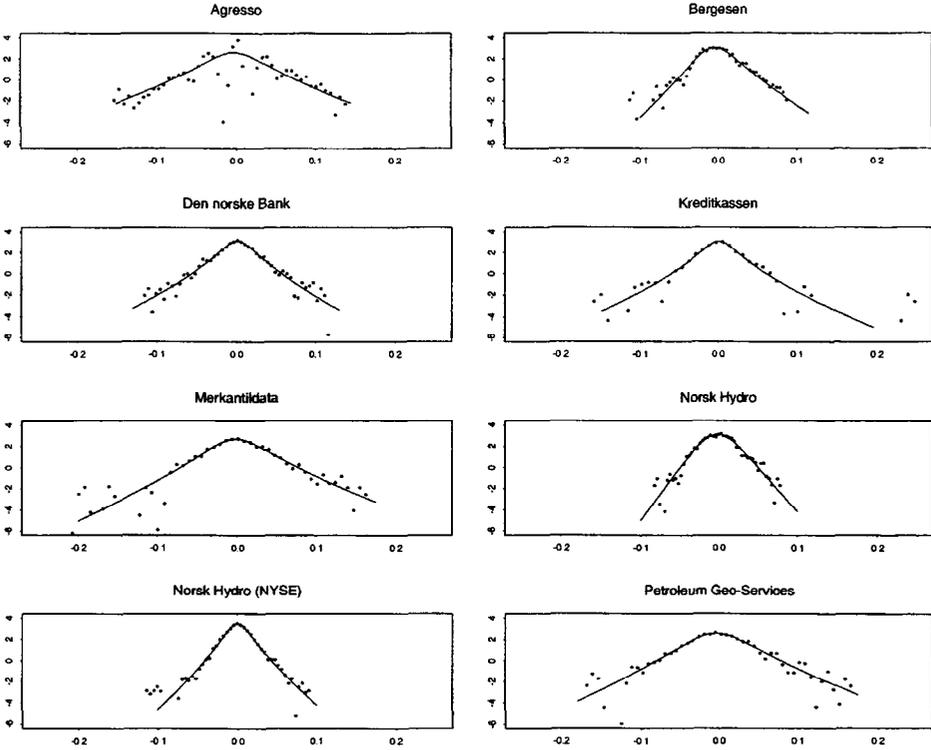


FIGURE 2. The fitted normal inverse Gaussian densities (solid lines) against non-parametric estimates (dots). Logarithmic scale on vertical axis.

three models was performed on a time horizon ranging from 1 day to nearly a month. The normal distribution gave consistently too optimistic Value-at-Risk estimates for all stocks in question compared to the non-parametric estimate. This is expected in view of the poor fitting of the tails. The normal inverse Gaussian distribution, on the other hand, seems to be closer to the non-parametric Value-at-Risk estimate. We also calculated the standard error based on bootstrapping techniques. The estimates from the non-parametric model showed high errors, both relative to its estimated Value-at-Risk and to the two other statistical models. This clearly indicates the uncertainty (and danger!) connected to Value-at-Risk estimation for percentiles far out in the tails.

By going through the same analysis on data divided into periods about half a year each, we found a huge upward bias in the estimate of α for the seven Norwegian stocks.

QUANTIFICATION OF RISK IN NORWEGIAN STOCKS

Stock	VaR 1%	1 day	5 days	10 days
		$\widehat{\text{VaR}}$	$\widehat{\text{VaR}}$	$\widehat{\text{VaR}}$
Agresso	nig	0.113 (0.013)	0.226 (0.021)	0.311 (0.024)
	normal	0.095 (0.005)	0.208 (0.014)	0.290 (0.025)
	non-par.	0.104 (0.040)	0.211 (0.041)	0.291 (0.041)
Bergesen	nig	0.056 (0.004)	0.119 (0.007)	0.168 (0.009)
	normal	0.054 (0.003)	0.116 (0.009)	0.160 (0.014)
	non-par.	0.060 (0.023)	0.124 (0.023)	0.173 (0.123)
Den Norske Bank	nig	0.074 (0.005)	0.147 (0.009)	0.200 (0.011)
	normal	0.060 (0.003)	0.133 (0.008)	0.188 (0.014)
	non-par.	0.073 (0.022)	0.139 (0.023)	0.192 (0.023)
Kredittkassen	nig	0.081 (0.008)	0.160 (0.014)	0.219 (0.017)
	normal	0.067 (0.004)	0.152 (0.012)	0.216 (0.020)
	non-par.	0.089 (0.029)	0.161 (0.030)	0.219 (0.032)
Merkantildata	nig	0.093 (0.009)	0.192 (0.013)	0.268 (0.017)
	normal	0.087 (0.005)	0.197 (0.014)	0.283 (0.023)
	non-par.	0.081 (0.037)	0.204 (0.037)	0.277 (0.037)
Norsk Hydro	nig	0.049 (0.004)	0.106 (0.005)	0.149 (0.007)
	normal	0.047 (0.002)	0.103 (0.007)	0.144 (0.012)
	non-par.	0.045 (0.023)	0.106 (0.022)	0.140 (0.022)
Norsk Hydro NYSE	nig	0.044 (0.002)	0.089 (0.003)	0.124 (0.004)
	normal	0.039 (0.001)	0.087 (0.003)	0.123 (0.005)
	non-par.	0.043 (0.014)	0.093 (0.015)	0.127 (0.016)
Petroleum Geo services	nig	0.101 (0.007)	0.209 (0.012)	0.291 (0.015)
	normal	0.091 (0.005)	0.199 (0.014)	0.277 (0.024)
	non-par.	0.105 (0.041)	0.210 (0.039)	0.294 (0.039)

TABLE 2. Estimates of 1% Value-at-Risk from fitted inverse normal Gaussian and ordinary Gaussian models and from a non-parametric procedure; estimated standard error in parenthesis.

This led to significantly lower Value-at-Risk estimates than the corresponding Value-at-Risk based on the complete dataset. Interestingly, the same phenomenon was *not* observed for the New York stock. Here, as one would a priori expect, Value-at-Risk based on subsets of data enveloped the estimate from the full data series.

REFERENCES

- [1] M. Abramowitz and I.A. Stegun (1972): Handbook of Mathematical Functions. Dover Publ. Inc. New York.
- [2] O. E. Barndorff-Nielsen, Exponentially decreasing distributions for the logarithm of particle size, *Proc. Roy. Soc. London A* **353** (1997), 401-419.

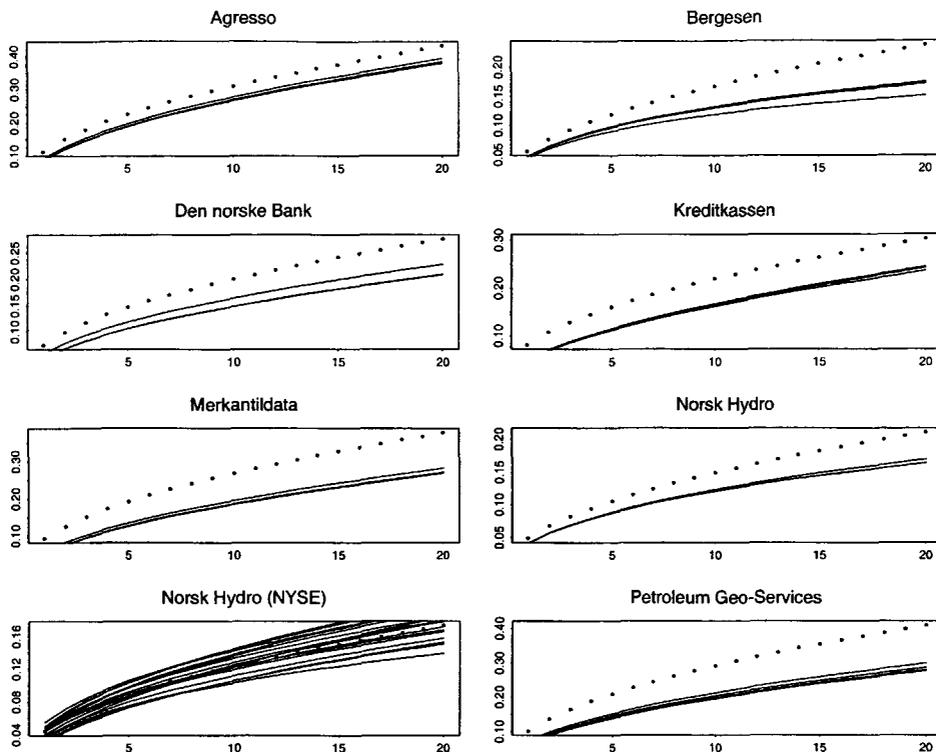


FIGURE 3. Estimated 1% Value-at-Risk based on half-year data (solid lines) and all data (dotted lines).

- [3] O. E. Barndorff-Nielsen, Processes of Normal inverse Gaussian type, *Finance and Stochastics* **2** (1998), 41-68.
- [4] O. E. Barndorff-Nielsen and O. Halgreen, Infinite divisibility of the hyperbolic and generalized inverse Gaussian distributions, *Z. Wahrscheinlichkeitstheorie Verw. Geb.* **38** (1977), 309-312.
- [5] P. Blæsild and M. K. Sørensen, Research Report No. 248. Department of Theoretical Statistics, University of Aarhus (1992).
- [6] E. Eberlein and U. Keller, Hyperbolic distributions in finance, *Bernoulli* **1**(3) (1995), 281-299.
- [7] E. Eberlein, U. Keller and K. Prause, New insights into smile, mispricing and value at risk: the hyperbolic model, *J. Business*, **71**(3) (1998), 371-405.
- [8] Efron and Tibshriani, An introduction to the bootstrap. New York: Wiley, 1993.
- [9] P. Embrechts, C. Klüppelberg and Th. Mikosch, Modelling Extremal Events. For Insurance and Finance. Applications of Mathematics, 33. Springer-Verlag, Berlin, 1997.
- [10] P. E. Gill, W. Murray and M. H. Wright, Practical optimization. Academic Press, Inc., London-New York, 1981.