

STOCHASTIC MODELLING OF TEMPERATURE VARIATIONS WITH A VIEW TOWARDS WEATHER DERIVATIVES

FRED ESPEN BENTH AND JŪRATĖ ŠALTYTĖ-BENTH

ABSTRACT. We model the daily average temperature variations with a mean-reverting Ornstein-Uhlenbeck process driven by generalized hyperbolic Lévy process and having seasonal mean and volatility. It is empirically demonstrated that the proposed dynamics fits Norwegian temperature data quite successfully, and in particular explains the seasonality, heavy tails and skewness observed in the data. The stability of mean-reversion and the question of fractionality of the temperature data are discussed. We apply our model to derive explicit prices for some standardized futures contracts based on temperature indices and options on these traded on the Chicago Mercantile Exchange (CME).

1. INTRODUCTION

Financial derivatives contracts based on weather conditions have gained increasing popularity over the recent years as a tool to manage risk exposure towards unfortunate weather events. Standardized futures contracts written on temperature indices have been traded on the Chicago Mercantile Exchange (CME) since October 2003, together with European call and put options written on these futures. The temperature indices are based on measurement locations in the US and Europe. To understand the risk involved in such trading, confident pricing models of both futures and options are called for.

In this paper we will study the mean-reverting Ornstein-Uhlenbeck stochastic process proposed by Dornier and Querel [16], and suggest some extensions of this motivated from an empirical investigation of more than 13 years of daily average temperature measurements at 7 locations spread over Norway. Based on our suggested Ornstein-Uhlenbeck model, we will derive a pricing dynamics for the futures and options commonly traded in weather derivatives markets like the CME. Furthermore, we will evaluate other versions of the Ornstein-Uhlenbeck model (see Alaton, Djehiche and Stillberger [3], Campbell and Diebold [13] and Brody, Syroka and Zervos [12]) in view of the Norwegian temperature data.

Date: January 6, 2004.

Key words and phrases. Temperature modelling, stochastic processes, Lévy processes, mean-reversion, seasonality, fractionality, temperature futures and options.

We are grateful to Kjersti Ulseth and Neil Shephard for valuable and inspiring discussions. Jūratė Šaltytė-Benth acknowledges the financial support by the Norwegian Research Council under grant NFR: 155120/432.

The literature suggests different types of stochastic mean-reversion models for the time dynamics of temperature extending that of Dornier and Querel [16]. In this paper we propose an Ornstein-Uhlenbeck model with seasonal mean and volatility, where the residuals are generated by a Lévy process rather than a Brownian motion. In particular, we suggest to use the class of generalized hyperbolic Lévy processes, a flexible class of Lévy processes capturing the semi-heavy tails and skewness which we observe in some of the Norwegian temperature data. The density and characteristic function of the generalized hyperbolic distribution is explicitly known, making it a convenient class of distributions for empirical study and derivatives pricing. Furthermore, it has the normal distribution as a limiting case.

The Ornstein-Uhlenbeck process of Dornier and Querel [16] models the temperature fluctuations as a regression between daily deseasonalized temperatures. We estimate the regression parameter, and evaluate the stability of this for the Norwegian temperature observations. The question of stability is of importance for derivatives pricing. Furthermore, we propose to model the temperature variance (or volatility) as an empirically based function estimated from the daily observed variances. This model for the variance function explains quite successfully the observed seasonality in the empirical correlograms for the squared residuals. After dividing by the square-root of this variance function from the residuals, we find that for some locations the normal distribution is rejected as a model, and other, more heavy-tailed and skewed distributions are called for. Hence, Lévy processes enter as a natural modelling tool for the dynamics of these residuals. In our empirical studies, we also address the question of spatial correlation between measurement locations and the fractional behaviour of temperature data.

Since our stochastic dynamics for temperature variations is a continuous-time diffusion model, it lends itself to the arbitrage theory for pricing of derivatives. In the classical Black, Merton and Scholes framework (see Black and Scholes [7] and Merton [20]), derivatives are assumed to be perfectly replicable, a natural condition for options written on tradeable assets like stocks. However, for futures written on temperature indices, we can not any longer base our valuation on hedging principles, since the underlying can not be traded. Thus, in modelling the pricing dynamics of futures on temperature, it is natural to include a parameter measuring the “market price of risk”. We will do so using the Esscher transform, a convenient approach in the context of Lévy processes. For futures contracts based on cumulative temperature we will be able to derive an explicit pricing dynamics expressed in terms of the market price of risk, temperature volatility and the characteristic function of the Lévy process. Furthermore, in the special case of Brownian motion, we can derive an expression for price of a call option written on such temperature futures. Numerical procedures are outlined for the general case.

The paper is organized as follows: in Section 2 we present our stochastic model for the daily temperature variations, together with a discussion of existing models which has been used for pricing weather derivatives. Our proposed Ornstein-Uhlenbeck model is fitted to Norwegian temperature data in Section 3, where we discuss the potentials and shortcomings of the model compared to other Ornstein-Uhlenbeck dynamics. In particular, we emphasize the seasonality of volatility, heavy tails of residuals and the question of fractionality in the

temperature data. In Section 4 we discuss pricing of futures and options on these based on our dynamical model for temperature. The paper ends with concluding remarks and a discussion of topics for further research in Section 5.

2. STOCHASTIC DYNAMICS OF TEMPERATURE VARIATIONS

In this Section we discuss different stochastic models for temperature variations. We suggest an Ornstein-Uhlenbeck process driven by Lévy noise to model temperature fluctuations, but also present in detail other models proposed in the literature.

Let (Ω, \mathcal{F}, P) be a complete probability space equipped with a filtration $\{\mathcal{F}_t\}_{t \geq 0}$ satisfying the *usual hypotheses* (see e.g. Karatzas and Shreve [19]). Introduce a Brownian motion $B(t)$ and an independent Lévy process $L(t)$. The Lévy process is assumed to be a pure-jump square-integrable process, and we choose to work with the version of $L(t)$ being right-continuous and having left-hand limits (the so-called càdlàg version). The Lévy measure of $L(t)$ is a σ -finite measure on the Borel sets of $\mathbb{R} \setminus \{0\}$ denoted by $\ell(dz)$ and satisfies the integrability condition

$$\int_{\mathbb{R} \setminus \{0\}} 1 \wedge z^2 \ell(dz) < \infty.$$

Here, $a \wedge b$ denotes the minimum of the two numbers a and b . The Lévy-Kintchine representation of $L(t)$ is

$$L(t) = \int_{\mathbb{R} \setminus \{0\}} z \tilde{N}(dt, dz),$$

with $N(dt, dz)$ being the homogeneous Poisson random measure associated to $L(t)$ and $\tilde{N}(dt, dz) := N(dt, dz) - \ell(dz)dt$ its compensated (Poisson) random measure. Throughout this paper we denote $T(t)$ the temperature at time $0 \leq t < \infty$.

2.1. Discussion of existing mean-reversion models. Dornier and Querel [16] and Alaton *et al.* [3] suggest the following Ornstein-Uhlenbeck dynamics for temperature variations:

$$(2.1) \quad dT(t) = ds(t) + \kappa(T(t) - s(t))dt + \sigma(t)dB(t),$$

where $s(t) = A + Bt + C \sin(\omega t + \phi)$ describes the mean seasonal variation (frequently referred to as the annual cycle or seasonality of the temperature) and the constant κ is the speed the temperature reverts to its mean. The volatility $\sigma(t)$ is assumed to be a measurable and bounded function. The model (2.1) regresses the *change* in deseasonalized temperature against deseasonalized temperature. As pointed out by Dornier and Querel [16], this model will tend towards a historical mean $s(t)$, which is not the case if the term $ds(t)$ on the right-hand side of (2.1) is left out, a highly undesirable property according to Dornier and Querel [16]. Even though Dornier and Querel [16] allow for a varying volatility function $\sigma(t)$, they assume it constant in their analysis of 20 years of daily average temperature data recorded in Chicago (US).

Alaton *et al.* [3] model $\sigma(t)$ as a piecewise constant function representing a monthly variation in volatility. Considering a data series covering 40 years of daily average temperatures from Bromma (nearby Stockholm, Sweden), Alaton *et al.* [3] fit the Ornstein-Uhlenbeck

model. They observe that the quadratic variation $\sigma^2(t)$ is nearly constant over each month in the data set, validating their choice of volatility function. Their argument for using a Wiener process as the driving noise in the Ornstein-Uhlenbeck process comes from the observation that the temperature differences are close to normally distributed. However, a statistical test for normality is not provided, and the authors admit that the empirical frequency of small temperature differences are higher than predicted by the fitted normal distribution. For the Norwegian temperature data, the normal hypothesis is rejected in several locations, and thus other, non-normal models are called for.

In neither of the two papers mentioned above there is a study of the possible time-dependencies in the residuals observed from the regression model. This time risk is the subject of investigation in Brody *et al.* [12], where a *fractional* Ornstein-Uhlenbeck process is suggested¹:

$$(2.2) \quad dT(t) = \kappa(T(t) - s(t)) dt + \sigma(t)dB^H(t),$$

with B^H being a *fractional* Brownian motion (see e.g. Shiryaev [22] for a definition). In (2.2) the change in temperature is regressed against previous day's *deseasonalized* temperature. However, following the arguments of Doernier and Querel [16], one should add the changes of seasonal variation $ds(t)$ on the right-hand side in order to have a consistent mean-reversion model.

Based on a data series of daily temperatures from central England recorded from 1772 up to 1999, Brody *et al.* [12] found clear signs of a fractional behaviour in the temperature fluctuations after the annual cycle $s(t)$ was removed. However, they do not perform the same fractional analysis for the residuals in their regression model. In fact, they do not use their data to fit the model (2.2) at all. It is not clear if the time-dependencies of the residuals (even after including a non-constant volatility $\sigma(t)$) will follow the characteristics of a fractional noise. As we shall see, a fractional Brownian dynamics does not seem to be an appropriate model when considering the Norwegian temperature data.

Campbell and Diebold [13] propose an autoregressive time-series to model temperature variations. Their model does not have any natural continuous-time analogue, but we include it here since Campbell and Diebold [13] report some very interesting results in their analysis of temperature data from the US. Let T_t evolve according to the autoregressive time series

$$(2.3) \quad T_t = m_t + s_t + \sum_{l=1}^L \rho_{t-l} T_{t-l} + \tilde{\epsilon}_t, \quad t = 1, 2, \dots,$$

where the trend m_t is linear and the seasonality s_t is modelled by a finite sum of sines and cosines. Note that we will use time t as a sub-index when considering time series, and reserve the notation $T(t), s(t), \dots$ etc. for the continuous-time models. Before presenting their choice of residual model, we want to comment on the structure of this temperature dynamics. It deviates significantly from the Ornstein-Uhlenbeck model (2.1) since the

¹In their model, Brody *et al.* [12] allow the speed of mean-reversion κ to vary with time. However, they do not discuss the modelling this any further and we find it most natural to leave it constant.

latter is only a first-order autoregression, whereas (2.3) regresses today's deseasonalized (and de-trended) temperature against the temperatures observed over the last L days. There is, however, one more major difference. If we put $L = 1$ in (2.3), we *do not* recover the Ornstein-Uhlenbeck models discussed above. In (2.1) we regress on daily temperatures *after* subtracting seasonality, while in (2.3) one is regressing *deseasonalized* temperature on previous day's *absolute* temperature. As we have already mentioned, Dornier and Querel [16] argue against this modelling perspective.

When analysing data from several cities in the US, Campbell and Diebold [13] find that the squared observed residuals $\tilde{\epsilon}_t^2$ have a very dominant seasonal variation over lags up to 800 days, while the residuals itself appear to be uncorrelated. To explain this phenomenon in terms of a statistical model, they suggest to use an autoregressive conditional heteroscedastic (ARCH) dynamics² $\tilde{\epsilon}_t = \sigma_t \epsilon_t$, where ϵ_t are iid and

$$(2.4) \quad \sigma_t^2 = \sum_{q=1}^Q \{ \gamma_{c,q} \cos(2\pi q d(t)/365) + \gamma_{s,q} \sin(2\pi q d(t)/365) \} + \sum_{r=1}^R \alpha_r \tilde{\epsilon}_{t-r}^2.$$

Here, $d(t)$ is a repeating step function that cycles through 1 to 365. As we are going to see, the Norwegian data possesses similar seasonality features as observed by Campbell and Diebold [13]. However, we will suggest a much simpler model than (2.4) to explain this variation.

2.2. A Lévy-based Ornstein-Uhlenbeck model. We propose the following generalization of the Ornstein-Uhlenbeck model (2.1) for the time evolution of temperatures:

$$(2.5) \quad dT(t) = ds(t) + \kappa (T(t) - s(t)) dt + \sigma(t) dL(t).$$

The only difference from (2.1) is the inclusion of a Lévy noise $L(t)$ rather than Brownian motion. We suggest to use a Lévy process with marginals following the class of generalized hyperbolic distributions. This is a very flexible family of distributions which can model skewness and (semi-)heavy tails. Furthermore, its density and characteristic (moment generating) functions are explicitly known.

The generalized hyperbolic distributions is a family of infinitely divisible distributions with density function

$$(2.6) \quad f_{\text{gh}}(x; \lambda, \mu, \alpha, \beta, \delta) = c (\delta^2 + (x - \mu)^2)^{(\lambda - \frac{1}{2})/2} \exp(\beta(x - \mu)) \times \\ K_{\lambda - \frac{1}{2}}(\alpha \sqrt{\delta^2 + (x - \mu)^2}).$$

Here K_s is the modified Bessel function of the third kind with index s (see e.g. [1], Section 9.6), and the normalizing constant c is given as

$$c = \frac{(\alpha^2 - \beta^2)^{\lambda/2}}{\sqrt{2\pi} \alpha^{\lambda - \frac{1}{2}} \delta^\lambda K_\lambda(\delta \sqrt{\alpha^2 - \beta^2})}.$$

The parameter α controls the steepness (or the fatness of the tails) of the distribution, μ the location of the distribution, β the skewness and δ is the scaling. The distribution is

²We believe there is a small misprint in Eq. (2) and (2c) in [13] in the presentation of the ARCH model

symmetric when $\beta = 0$. Finally, the parameter λ is identifying the sub-family within the generalized hyperbolic class. Two special cases of the generalized hyperbolic distribution which have been much studied in a financial context are the hyperbolic distribution with $\lambda = 1$, and normal inverse Gaussian distribution with $\lambda = -0.5$. The hyperbolic distribution was suggested by Eberlein and Keller [17] as a flexible model for financial logreturns, while Rydberg [21] and Barndorff-Nielsen [8] considered the normal inverse Gaussian (NIG) distribution. The generalized hyperbolic family has the normal, Student T and the Cauchy distributions as limiting cases.

The family of generalized hyperbolic distributions admits an explicit moment generating function. If a random variable X is distributed according to a generalized hyperbolic distribution, the moment generating function is,

$$(2.7) \quad \mathbb{E}[e^{uX}] = e^{\mu u} \left(\frac{\alpha^2 - \beta^2}{\alpha^2 - (\beta + u)^2} \right)^{\lambda/2} \frac{K_\lambda(\delta \sqrt{\alpha^2 - (\beta + u)^2})}{K_\lambda(\delta \sqrt{\alpha^2 - \beta^2})},$$

whenever $|\beta + u| < \alpha$. Hence, this family of distributions have finite moments of all orders.

We call $L(t)$ a *generalized hyperbolic Lévy process* if $L(t)$ is a Lévy process with the marginals $L(1)$ being distributed according to the generalized hyperbolic family. In this case the Lévy measure is given explicitly by

$$(2.8) \quad \ell_{\text{GH}}(dz) = |z|^{-1} e^{\beta z} \left\{ \frac{1}{\pi^2} \int_0^\infty \frac{\exp(-\sqrt{2y + \alpha^2}|z|)}{J_\lambda^2(\delta\sqrt{2y}) + Y_\lambda^2(\delta\sqrt{2y})} \frac{dy}{y} + \lambda e^{-\alpha|z|} \right\} dz,$$

when $\lambda \geq 0$, and

$$(2.9) \quad \ell_{\text{GH}}(dz) = |z|^{-1} e^{\beta z} \frac{1}{\pi^2} \int_0^\infty \frac{\exp(-\sqrt{2y + \alpha^2}|z|)}{J_{-\lambda}^2(\delta\sqrt{2y}) + Y_{-\lambda}^2(\delta\sqrt{2y})} \frac{dy}{y} dz,$$

when $\lambda < 0$. Here J_λ and Y_λ are the Bessel functions of the first and second kind, respectively, with index λ (see e.g. [1], Section 9.1.). The generalized hyperbolic Lévy processes $L(t)$ are pure-jump processes with paths of infinite variation. For more theory on the class of generalized hyperbolic distributions and the associated Lévy processes we refer the interested reader to Barndorff-Nielsen and Shephard [10].

A straightforward application of the Itô Formula for semimartingales (see e.g. Ikeda and Watanabe [18]), leads to an explicit solution of (2.5):

$$(2.10) \quad T(t) = s(t) + (T(0) - s(0)) e^{\kappa t} + \int_0^t \sigma(u) e^{\kappa(u-t)} dL(u).$$

Furthermore, the cumulative temperature over a time interval $[\tau_1, \tau_2]$ is easily calculated:

Lemma 2.1. *If the temperature $T(t)$ follows (2.5), the cumulative temperature over the time interval $[\tau_1, \tau_2]$ is explicitly given by*

$$(2.11) \quad \int_{\tau_1}^{\tau_2} T(t) dt = \int_{\tau_1}^{\tau_2} s(t) dt + \kappa^{-1} (T(0) - s(0)) (e^{\kappa\tau_2} - e^{\kappa\tau_1}) \\ + \int_0^{\tau_2} \sigma(t) \kappa^{-1} \{ e^{\kappa(\tau_2-t)} - \mathbf{1}_{[0, \tau_1]}(t) e^{\kappa(\tau_1-t)} - \mathbf{1}_{[\tau_1, \tau_2]}(t) \} dL(t).$$

Proof. Let $\tilde{T}(t) = T(t) - s(t)$ be the deseasonalized temperature. From (2.5) we find

$$\tilde{T}(\tau_2) = \tilde{T}(\tau_1) + \kappa \int_{\tau_1}^{\tau_2} \tilde{T}(t) dt + \int_{\tau_1}^{\tau_2} \sigma(t) dL(t).$$

Combining this with the explicit dynamics of $T(t)$ in (2.10) yield the Lemma. \square

The explicit form in (2.11) is useful in connection with derivatives written on the cumulative temperature as underlying (see Section 4).

We end this Section by deriving a time-discrete version of (2.5): letting time be measured in days, say, we have the model

$$\Delta T(t) = \Delta s(t) + \kappa(T(t-1) - s(t-1))\Delta t + \sigma(t-1)\Delta L(t)$$

with $\Delta t = 1$ and $\Delta Y(t) = Y(t) - Y(t-1)$, where $Y(t)$ is short-hand for the processes $T(t)$, $L(t)$ or $s(t)$. After reorganizing, we reach the following time-series analogue of (2.5)

$$(2.12) \quad T_t - s_t = (1 + \kappa)(T_{t-1} - s_{t-1}) + \sigma_t \epsilon_t, \quad t = 1, 2, \dots$$

This will be our basic discrete-time model when we analyse the Norwegian temperature data in the next Section.

3. ANALYSIS OF NORWEGIAN TEMPERATURE DATA

We analyse a dataset of daily mean temperatures (measured in centigrades) observed in 7 cities in Norway over a period ranging from 1990 01 01 to 2003 04 08, resulting in 7 data series of 4846 observations each ³. The cities we consider are Alta, Bergen, Kristiansand, Oslo, Stavanger, Tromsø and Trondheim, all being located along the coast of Norway. Alta and Tromsø are the northern-most cities, while Trondheim is located in the middle of Norway. Bergen and Stavanger are situated at the western coast of South-Norway, while Kristiansand is the southern-most city. Oslo is placed in the eastern part of South-Norway.

In Table 1 we list descriptive statistics characteristics associated to the 7 measurement stations (cities). In addition, we display the values of the χ^2 -statistics of Pearson's criteria of goodness-of-fit with the corresponding P -values. The mean, median, minimum and maximum values of daily mean temperatures differ from city to city, but this is explainable by quite distinct and wide-spread geographical locations. Standard deviations (std) are not big, but also differ among the cities. The shape of the empirical distributions is not symmetrical (values of skewness are different from zero) and has negative kurtosis; in some cases we observe a multimodal pattern, being a clear sign of strong seasonality. Indications about the non-normal shape of the empirical distributions are confirmed by the last two rows in Table 1: the values of χ^2 -statistics are significant at the 1% level for all 7 cities.

Our proposed model (2.12) of the daily average temperature can be written as the following additive time series

$$(3.1) \quad T_t = s_t + c_t + \tilde{\epsilon}_t, \quad t = 0, 1, 2, \dots$$

Here, T_t is the average temperature on day t , s_t the seasonal component, c_t the cyclical component, and $\tilde{\epsilon}_t$ the noise. In the rest of this Section we will discuss our choices for

³February 29 was removed from the sample in each leap year to have years of equal size

the different components and fit the model to the observed temperature variations at the different locations. The statistical analysis that follows revealed that the seven cities could be collected into two groups, each with their typical pattern: Alta, Kristiansand, Stavanger, Tromsø and Trondheim constitute one group, and Bergen and Oslo the other. Hence, we shall frequently illustrate our findings using Alta and Bergen as typical representatives for the two groups.

TABLE 1. Descriptive statistics for daily average temperature

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
Mean	2.01	8.29	7.43	6.75	8.02	3.21	6.07
Median	2.20	7.85	7.20	6.20	7.70	2.80	5.80
Mode	3.70	multiple	3.85	multiple	4.60	3.10	multiple
Std	8.60	5.77	6.76	8.09	5.75	6.51	7.23
Min.	-30.80	-9.60	-16.50	-18.15	-10.60	-13.05	-22.45
Max.	25.80	25.10	24.60	24.00	24.15	20.50	26.45
Skewness	-0.23	0.07	-0.20	-0.10	-0.04	0.03	-0.20
Kurtosis	-0.36	-0.59	-0.46	-0.76	-0.39	-0.70	-0.16
χ^2	217.55	240.57	390.27	517.60	227.43	347.97	176.63
P -value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

3.1. Trend. Alaton *et al.* [3], and Campbell and Diebold [13] observe that over several decades there has been a clear increase of daily mean temperatures in Sweden and USA, respectively. Such an increase in average temperature can be explained by global warming, green-house effect, urbanization, or other phenomena. In our data analysis we did not find any significant linear trend: the value of R^2 associated with the regression line was below 0.21% for all 7 stations. The most likely reason for this is that our time series are too short: they contain temperature data recorded for less than 14 years, while the datasets of Alaton *et al.* [3] and Campbell and Diebold [13] range over approximately 40 years. Over the considered time span, we conclude that there is no typical trend in the daily mean temperatures in Norway, and assume this equal to zero in the rest of the paper.

Since our view is towards modelling daily temperature variations for a short time horizon (for e.g. up to one year) as a basis for analyzing weather derivatives, the trend will not give any significant influence (see the results of Alaton *et al.* [3], which give a very small contribution over the short time intervals we have in mind). Also, in the long run it is questionable whether there will be a steady increase in temperature given by a linear trend.

3.2. Seasonality. It is natural to expect that temperature follows a certain seasonal pattern: we always experience lower temperatures in the winter and higher in the summer. We model this seasonal dependency (the annual cycle of temperature) with a simple cosine

function

$$(3.2) \quad s_t = a_0 + a_1 \cos\left(\frac{2\pi}{365}(t - t_0)\right),$$

where a_0 and a_1 are constants describing the average level and amplitude of the mean temperature, respectively, and t_0 defines a phase angle. In Table 2 we report the different estimated parameters⁴ for each station. Figure 1 displays the average daily temperature in Alta and Bergen along with their estimated seasonal component s_t .

TABLE 2. Estimated parameters for the seasonality function

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
a_0	2.17	8.41	7.56	6.90	8.13	3.33	6.20
a_1	-10.21	-6.91	-8.28	-10.28	-6.86	-7.73	-8.45
t_0	20.79	21.64	20.08	15.31	24.43	22.99	17.68

3.3. Cyclicity and Regression Analysis. From (2.12), we model the cyclical component by regressing today's deseasonalized temperature against the deseasonalized temperature recorded yesterday, resulting in:

$$(3.3) \quad c_t = \alpha(T_{t-1} - s_{t-1}),$$

where $\alpha = 1 + \kappa$. Hence,

$$(3.4) \quad X_t = \alpha X_{t-1} + \tilde{\epsilon}_t,$$

using the notation $X_t := T_t - s_t$. The values of the fitted slopes α and the corresponding R^2 are given in Table 3 for all 7 cities. We see that the values of R^2 are all reasonably good. Campbell and Diebold [13] fit their model (2.3) to ten U.S. cities using 25 autoregressive

TABLE 3. Values of α and R^2 in regression analysis

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
α	0.77	0.81	0.76	0.80	0.80	0.81	0.79
R^2	59.0%	65.6%	57.1%	63.9%	63.6%	66.3%	63.0%

lags ($L = 25$) and 3 sine and cosine terms in s_t , and reach a fit with an R^2 above 90%. We tried the same model for the Norwegian data where we only reached an R^2 of 40% (this was done for Alta). As we see from Table 3, the R^2 for our considerably simpler model (2.12) with just one cosine function in s_t is 59% for Alta. Thus, we conclude that our regression model performs much better for Norwegian data than (2.3). We believe that it is much better to do regression on the *deseasonalized* temperatures, a claim that is supported by looking at the scatter plots in Fig. 2 for the two regression models for Alta. We remark

⁴We used `nlinfit` in Matlab to perform this estimation.

in passing that, as expected after deseasonalizing, the intercepts in the regression were not significantly different from zero. This explains why we have not included them in our regression model (3.4).

In Fig. 3 we plot the residuals and squared residuals for Alta and Bergen over the observed time period. In almost all cities we find a clear persistent variation in the variance of noise, a sign of seasonal heteroscedasticity. For Alta this is very apparent, especially when we consider the squared residuals. For Bergen, however, the picture is not as clear, although signs of seasonal heteroscedasticity in noise can be observed here as well. The findings suggest that one should include a deterministic seasonality function when modelling variance as a certain function of time.

Campbell and Diebold [13] observed that the squared residuals from their regression model showed strong seasonality effects in its autocorrelation function (ACF). They suggest an ARCH model for the residual dynamics where $\tilde{\epsilon}_t = \sigma_t \epsilon_t$ and σ_t^2 follows (2.4). We propose a simpler and more intuitive model for σ_t as a seasonally varying deterministic function estimated from the empirical variance. Our findings, that we discuss below, are in line with those reported in Campbell and Diebold [13], and we see no reason to introduce an ARCH dynamics in our case. We remark that Campbell and Diebold [13] do not investigate the claimed normality of $\epsilon_t := \tilde{\epsilon}_t/\sigma_t$ (see below for a study of this for the Norwegian data).

The empirical densities of the residuals of Alta and Bergen are plotted in Fig. 4, and Table 4 displays the descriptive statistics characteristics for all cities. The values of χ^2 are again all significant at the level 1%, rejecting the normal hypothesis. This is evident from Fig. 4, where we have plotted the density of the maximum likelihood fitted normal distribution together with the empirical density. We note that the normal distribution overestimates the probability for small variations and does not capture the tail behaviour of the empirical density. The latter effect is highlighted by using a logarithmic scale on the frequency axis in the plots. We observe that the signs of non-normality are much more evident for Alta than Bergen. Noteworthy is that the kurtosis of the residuals for all 7 cities is now positive, while we recall from Table 1 that the daily average temperatures had negative kurtosis. Also, the variability in residuals is much smaller than in average temperature itself.

Modelling the residuals $\tilde{\epsilon}_t$ by a normal distribution with constant variance (like in Dornier and Querel [16]) will lead to a significant underestimation of large variations in the residuals. In the next Subsection we will see that a seasonally varying variance in conjunction with a heavy-tailed distribution explain the dynamics of the residuals $\tilde{\epsilon}_t$ reasonably well. We will argue for a simpler model than the one found in Campbell and Diebold [13], but more sophisticated than the one in Alaton *et al.* [3].

In Fig. 5 we have plotted the estimated ACF for the residuals and their squares for Alta and Bergen. The broken lines in the figure are the estimated 95% confidence intervals⁵ for having white noise residuals. We see that the autocorrelations for the residuals are roughly within this band for nearly all lags, with the exception of lags 1 and 2. For all cities we observed consistently that the ACF for lag 1 was positive, while ACF for lag 2 was

⁵The confidence band is given automatically by Splus

TABLE 4. Descriptive statistics for residuals after regression

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
Mean	0.00	0.00	-0.00	0.00	0.00	0.00	-0.00
Median	-0.07	-0.09	-0.05	-0.09	-0.08	-0.01	-0.08
Mode	no mode	no mode	no mode	no mode	no mode	no mode	no mode
Std	2.92	1.79	2.20	1.79	1.85	2.04	2.46
Min.	-12.98	-6.41	-11.10	-6.37	-7.83	-6.47	-15.55
Max.	15.97	7.15	8.47	7.21	7.44	7.83	10.30
Skewness	0.13	0.20	0.07	0.19	0.16	0.18	-0.01
Kurtosis	1.53	0.22	0.73	0.21	0.54	0.10	0.82
χ^2	341.64	63.18	102.79	56.80	111.12	80.54	77.88
P -value	0.0000	0.0004	0.0000	0.0022	0.0000	0.0000	0.0000

negative. For both lags the 95% confidence band is violated, so these lags are significantly different from zero. Unfortunately, we are not able to explain this in our chosen model.

The ACF in Fig. 5 for the squared residuals, on the other hand, reveals a clear seasonal pattern for both cities. The estimated 95% confidence band for white noise is violated for basically all lags, and the wavy form of the ACF is a clear sign of seasonal heteroscedasticity in the residuals. Admittedly, the picture is more clear for Alta than Bergen.

3.3.1. Stability Analysis of the Regression Model. It is of interest to understand how stable the regression parameter α is over time, and whether or not it varies with season. From Alaton *et al.* [3] we know that certain derivatives based on temperature are highly dependent on the mean-reversion parameter, and the variability of this will thus have a significant influence on the derivatives prices. To investigate this, we looked for possible yearly or seasonal variations in the estimates of α . Firstly, we estimated the regression parameter based on data from one year only, repeating this for every year until the end of the data sets. Secondly, we investigated the seasonal structure in α by estimating the regression parameter for each month individually, that is, we considered data for January only, then for February, and so on. In Tables 5-6 we report the obtained average α for each city together with the standard deviation for this. We also include the variation coefficient V (the std relative to the average) in percent, to illustrate the normal variations of α around its mean. Table 5 contains the figures for the yearly study, while the numbers for monthly study are displayed in Table 6. The variation of α is rather small over the years, with an exception of Alta which has a variation close to 6% of the average. The monthly variations of α are bigger for all cities (with Tromsø as the exception). Kristiansand is most prominent, with more than 9% variation over the year. From the investigations we could not observe any clear seasonal pattern for α in any of the cities, nor could we see any specific pattern over the years ⁶.

⁶Admittedly, this can not be seen from the tables. The estimated numbers can be obtained from the authors by request

TABLE 5. Values of average α estimated for each year together with the std and variation coefficient

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
average α	0.76	0.81	0.75	0.80	0.79	0.79	0.78
std	0.045	0.019	0.017	0.018	0.027	0.020	0.028
V	5.94%	2.41%	2.26%	2.29%	3.46%	2.58%	3.55%

TABLE 6. Values of average α estimated for each month together with the std and variation coefficient

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
average α	0.75	0.80	0.72	0.78	0.79	0.79	0.79
std	0.037	0.021	0.066	0.042	0.027	0.020	0.014
V	4.94%	2.57%	9.12%	5.37%	3.46%	2.58%	1.82%

Our conclusion is that the regression parameter α is modestly stable over years and months for most locations, and that no seasonality can be observed for this parameter. The latter implies that there is no reason to letting α be dependent on time.

3.4. Seasonality in the residuals $\tilde{\epsilon}_t$. Next we discuss modelling of the residuals $\tilde{\epsilon}_t$. Inspecting the time series of $\tilde{\epsilon}_t$ for Alta and Bergen (see Fig. 3), one observes a seasonality in the residuals with bigger variation in the winter season than in the summer. After applying Bartlett's test (see, e.g. Aivazian and Mkhitarian [2]) we reach the same conclusion: daily mean temperatures are heteroscedastic since the values of Bartlett's statistics are significant for all 7 cities at the level 1%. Hence, we propose a multiplicative time series model for the residuals $\tilde{\epsilon}_t$ given by

$$(3.5) \quad \tilde{\epsilon}_t = \sigma_t \epsilon_t,$$

where σ_t is a deterministic function and ϵ_t has zero mean and standard deviation equal to one. Since $\sigma_t^2 = E[\tilde{\epsilon}_t^2]$, the function σ_t^2 will describe the seasonal variation, of the residual's variance.

We describe how to fit σ_t^2 to the observed residuals $\tilde{\epsilon}_t = (T_t - s_t) - \alpha(T_{t-1} - s_{t-1})$ which we obtained from the regression described in the Subsection above: first, we organize the observed residuals into 365 groups, one for each day of the year. Finding the average of the squared residuals in each group, we obtain an estimate for the expected daily squared residual. This estimate is based on 13 observations only⁷, and hence the estimates become very noisy. We suggest to smoothen the daily estimates by taking the logarithm and then use a three-day moving-average. After exponentiating, we obtain a fitted (and smoothed) σ_t^2 as graphed in Fig. 6 for Alta and Bergen. The smoothing procedure decreases the scale

⁷Recall that we had approximately 13 years of observations, which means that each group will constitute of 13 residuals for each particular day of the year.

of the squared residuals and smoothens the variations, that is, we obtain a more regular estimate of σ_t^2 . Note that we assume σ_t^2 to be a periodic function such that $\sigma_t = \sigma_{t+k \cdot 365}$ for $t = 1, \dots, 365$ and $k = 1, 2, 3, \dots$

We note in passing that we tried to fit a Fourier series like the first sum in the ARCH-model (2.4) of Campbell and Diebold [13] (using $Q = 2$). The results were not particularly promising, and from the shape of our suggested σ_t^2 depicted in Fig. 6 one can see the reason why: one needs a large number of sines and cosines in order to explain this shape of the seasonality function.

The observed $\epsilon_t := \tilde{\epsilon}_t/\sigma_t$ and their squares are displayed in Fig. 7 for Alta and Bergen. The ACF for the observed ϵ_t and ϵ_t^2 of Alta and Bergen are plotted in Fig. 8. It is evident that the cyclical variation is removed completely for the squared residuals of Alta. For Bergen the cyclical pattern was not so obvious, but still one can see that introducing σ_t made some positive impact. From the plots we note that the ACF of the squared residuals are roughly within the 95% confidence interval for independence. However, it is worth observing that for all 7 cities the ACF for the squared residuals at lag 1 has a positive value in the interval from 0.05 to 0.09. Furthermore, for some cities it seems to fall exponentially down before it randomly fluctuates around 0.

We remark in passing that, as for $\tilde{\epsilon}_t$, we consistently have a negative correlation at lag 1 and positive correlation at lag 2 for the residuals ϵ_t , which are not explained by introducing a seasonality function σ_t .

3.5. The residuals ϵ_t modelled by the generalized hyperbolic distribution. In all the papers we have referred to above on temperature modelling (Doernier and Querel [16], Alaton *et al.* [3], Campbell and Diebold [13] and Brody *et al.* [12]), it is assumed that the residuals (after possibly modelling the variance) are iid normally distributed. In this Subsection we argue statistically that this is not always the case, and propose a flexible class of distributions for modelling of the residuals.

The values of the χ^2 -statistics and the corresponding P -values for the observed residuals ϵ_t are presented in Table 7. As we see from the results, we can *reject* the hypothesis that the residuals of Alta, Bergen and Oslo follow the normal distribution at a 1% significance level, while normality for Kristiansand is rejected at the 2% level. The residuals of Stavanger, Tromsø and Trondheim, on the other hand, can be confidently modelled by a normal distribution.

TABLE 7. Values of χ^2 statistics and corresponding P -values for residuals after dividing out the seasonal variation

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
χ^2	95.82	95.84	85.58	109.65	61.12	71.23	67.19
P -value	0.0017	0.0017	0.0134	0.0001	0.3996	0.1320	0.2169

Contrary to the conclusion in Alaton *et al.* [3] and Brody *et al.* [12], we can not always use a normal distribution to model the residuals for Norwegian temperature data. We propose

to apply the generalized hyperbolic family. The parameters of this family of distributions were fitted using maximum likelihood estimation⁸, where the results are reported for Alta, Bergen and Kristiansand in Table 8. We did not manage to fit the generalized hyperbolic distribution for Oslo data, but report in Table 8 the estimates we obtained after fixing the shape $\lambda = 30$ (which from Fig 11 seems to be a reasonable choice). The fitted generalized hyperbolic distribution together with the empirical distribution of Alta and Bergen are depicted in Fig. 10. As we see, the fit in the tails are improved compared with the normal distribution shown in Fig. 9. Furthermore, the skewness is now explained almost perfectly. Figure 11 shows the corresponding plots for Kristiansand and Oslo, where the same conclusions hold true.

TABLE 8. Estimates of the fitted generalized hyperbolic distribution

	μ	α	β	γ	λ
Alta	-0.35570	2.9633	0.32167	2.4899	1.3227
Bergen	-3.5227	11.876	4.3875	0.019906	48.861
Kristiansand	-1.4222	8.3041	1.4122	0.0012306	33.732
Oslo	-8.6747	26.282	18.735	6.6839	30*

3.6. Fractional analysis of temperature data. Based on measurements of daily temperature in central England from 1772 to 1999, Brody *et al.* [12] investigate the fractional behaviour of temperature fluctuations after the seasonality $s(t)$ has been removed. We now explain in more detail their approach (taken from Syroka and Toumi [23]) to estimate fractionality, and apply this to our data sets: divide the daily data (deseasonalized temperature) into N nonoverlapping bins of length K days, where N is greater than or equal to 10. Average the data within each bin to obtain \bar{X}_i for all $i = 1, \dots, N$. Define the statistics

$$f(K) := \sqrt{\frac{1}{N} \sum_{i=1}^N \bar{X}_i^2}.$$

According to Syroka and Toumi [23], uncorrelated data will give $\ln f(K) \sim -0.5 \ln K$, while the existence of a persistent temporal correlation can be concluded when $\ln f(K) \sim -(1 - H) \ln K$, for $H \in (0.5, 1)$. The number H is frequently called the *Hurst coefficient*, and was estimated in Brody *et al.* [12] to be $H = 0.61$ for the central England temperature data, a clear sign of fractionality. In Table 9 we report the corresponding numbers for the Norwegian data series together with the corresponding R^2 . As we can see, all the estimated Hurst coefficients have a high R^2 . Moreover, with the exception of Kristiansand, we find that the Hurst coefficients are greater than 0.5, indicating a persistent temporal dynamics. Bergen, Kristiansand and Oslo all have Hurst coefficients close to 0.5, which suggests temporally uncorrelated data, while the other cities show rather clear signs of fractionality.

⁸This is efficiently done using the programming language `0x`, see Doornik [15]

TABLE 9. Values of the Hurst coefficient (H) and R^2 in the fractional analysis of temperature

	Alta	Bergen	Kr. sand	Oslo	Stavanger	Tromsø	Trondheim
H	0.638	0.518	0.495	0.516	0.550	0.728	0.543
R^2	91.2%	91.0%	92.2%	91.0%	90.6%	89.8%	92.2%

When looking at model (2.2), one may argue that the fractionality is modelled in the wrong place. The analysis of fractionality in Brody *et al.* [12] concentrates on the temperature deviations away from the annual cycle, while the fractionality itself is modelled in the residuals from the regression of temperature changes. The authors do not provide any empirical study of the fractionality in the residuals.

If we consider carefully the autocorrelation function of the residuals ϵ_t obtained after dividing out the seasonal variation σ_t for the Norwegian cities (see Fig. 8 for Alta and Bergen), we do not observe any decay at a hyperbolic rate as predicted by the fractional Brownian motion. We claim that the fractionality observed for the temperature data in Norway can be better explained by a simple regression model for the deseasonalized temperature in conjunction with a time-dependent variance function σ_t , rather than introducing a fractional Brownian motion. Hence, a fractional model does not seem to be necessary for the Norwegian temperature data.

3.7. The Correlation between cities. In this Subsection we discuss the correlations in temperature between different cities. In table 10 the estimated correlations between the residuals after removing the seasonality s_t from the average temperature are presented. All cities are positively correlated, and we observe as expected that the correlation between cities far away from each other are rather small, while cities closely located are strongly correlated.

TABLE 10. Correlations between cities for the average temperature less seasonality

	Alta	Bergen	Kristiansand	Oslo	Stavanger	Tromsø	Trondheim
Alta	1.000	0.269	0.242	0.279	0.242	0.845	0.433
Bergen		1.000	0.740	0.750	0.910	0.301	0.718
Kristiansand			1.000	0.860	0.823	0.265	0.624
Oslo				1.000	0.779	0.297	0.690
Stavanger					1.000	0.271	0.700
Tromsø						1.000	0.485
Trondheim							1.000

If we estimate the correlations among the residuals ϵ_t obtained after dividing out σ_t in $\tilde{\epsilon}_t$, the conclusions remain the same, although with slightly different figures.

Understanding the correlations among cities are of major importance for an issuer of weather derivatives. It is not possible to diversify the risk by issuing temperature derivatives in several different locations. On the contrary, the risk is increased by the positive correlations observed. An investor with a portfolio of short positions in temperature options is likely to face liabilities from all the options, and not only some which can be outweighed by no exercise of other options. The investigation of the spatial correlation structure of temperature dynamics and its influence on derivatives written on temperature is a topic for future research.

4. DERIVATIVES ON TEMPERATURE

Chicago Mercantile Exchange (CME) offers standardized trading on futures and options written on temperature indices for several US and European cities⁹. The futures have the number of Heating-Degree Days (HDD)¹⁰ (or Cooling-Degree days (CDD)) over one month or one season for 15 US cities as underlying. For 5 European cities one can trade in futures written on the cumulative (average) temperature (CAT) over a season as well. The options written on these different futures contracts are plain vanilla European call and put options. We shall concentrate our considerations on the pricing of futures on CAT and options written on these, since they admit more or less explicit expressions.

We note that several authors have studied the problem of pricing options on temperature. In Alaton *et al.* [3] the fair value of a call option written on the number of HDD over a period is derived using a numerical approach, while Brody *et al.* [12] find the price of call options written on different combinations of HDD's as the solution of certain partial differential equations. Benth [4] generalizes the work of Brody *et al.* [12], and derives the time dynamics of temperature options based on a fractional dynamics.

Let us consider the price dynamics of futures written on CAT over a specified period $[\tau_1, \tau_2]$, $\tau_1 < \tau_2$ (the summer season, say). We suppose that the temperature dynamics follows (2.5) with $L(t)$ being a Lévy process (where $L(1)$ has finite exponential moments). Assuming a constant continuously compounding interest rate r , the futures price at time $t \leq \tau_1$ written on the CAT is defined as the \mathcal{F}_t -adapted stochastic process $F(t, \tau_1, \tau_2)$ satisfying

$$(4.1) \quad 0 = e^{-r(\tau_1-t)} \mathbb{E}_Q \left[\int_{\tau_1}^{\tau_2} T(t) dt - F(t, \tau_1, \tau_2) \mid \mathcal{F}_t \right].$$

Here Q is a risk-neutral probability. From the adaptedness of $F(t, \tau_1, \tau_2)$, we easily find the futures price to be

$$(4.2) \quad F(t, \tau_1, \tau_2) = \mathbb{E}_Q \left[\int_{\tau_1}^{\tau_2} T(t) dt \mid \mathcal{F}_t \right].$$

In order to derive a more explicit expression for the futures price, we need to specify the risk-neutral probability Q . Since temperature is not a storable commodity, futures

⁹See <http://www.cme.com/prd/wec/> for more information about this trading.

¹⁰The HDD index over the time interval $[\tau_1, \tau_2]$ is defined in a continuous-time setting as $\int_{\tau_1}^{\tau_2} \max(18 - T(t), 0) dt$. One considers 18°C to be the limit for switching *on* heating, thereby the name HDD.

contract can not be hedged at all and we have therefore an incomplete market. A risk-neutral probability is by definition a probability measure $Q\tilde{P}$ such that all *tradeable* assets in the market are martingales after discounting. Thus, *all* equivalent probabilities Q will become risk-neutral probabilities. Following the analysis in Benth and Saltyte-Benth [6], we specify a sub-family of probability measures using the so-called Esscher transform.

Assume $\theta(t)$ is a real-valued measurable and bounded function. Consider the stochastic process

$$(4.3) \quad Z^\theta(t) = \exp\left(\int_0^t \theta(s) dL(s) - \int_0^t \phi(\theta(s)) ds\right),$$

where $\phi(\lambda)$ is the moment generating function of $L(t)$, e.g. $\phi(\lambda) = \mathbb{E}[\exp(\lambda L(1))]$. The process $Z^\theta(t)$ is well-defined under natural exponential integrability conditions on the Lévy measure ℓ , which we assume to hold. We skip the exact reference to these conditions since they are rather technical and holds for all relevant models considered in this paper. The interested reader is referred to Benth and Saltyte-Benth [6] for a precise statement of the integrability conditions. Introduce the probability measure Q^θ defined by the Esscher transform:

$$Q^\theta(A) = \mathbb{E} [1_A Z^\theta(\tau_{\max})] ,$$

where 1_A is the indicator function and τ_{\max} is a fixed time horizon including the trading time for all relevant futures. This measure is obviously equivalent to P , and we denote the expectation under the probability Q^θ by $\mathbb{E}_\theta[\cdot]$. By using time-varying θ 's we have a flexible class of martingale measures Q^θ that we easily can fit to the observed forward curve. We call θ the “market price of risk”. Note in passing that if the Lévy process L is a Brownian motion, then the Esscher transform corresponds to a Girsanov change-of-measure, which in practice corresponds to a change of the drift in the temperature dynamics. We find the following dynamics for the futures contracts based on CAT's:

Theorem 4.1. *The futures price $F(t, \tau_1, \tau_2)$ at time $t \leq \tau_1$ written on a CAT over the interval $[\tau_1, \tau_2]$ is,*

$$(4.4) \quad F(t, \tau_1, \tau_2) = \int_{\tau_1}^{\tau_2} s(t) dt + \kappa^{-1} \left(e^{\kappa(\tau_2-t)} - e^{\kappa(\tau_1-t)} \right) \left(T(t) - s(t) \right) + \Theta(t, \tau_1, \tau_2),$$

where $\Theta(t, \tau_1, \tau_2)$ is given as a function of the market price of risk and volatility as

$$(4.5) \quad \Theta(t, \tau_1, \tau_2) = \kappa^{-1} \int_t^{\tau_2} \sigma(u) e^{\kappa(\tau_2-t)} \phi'(\theta(u)) du - \kappa^{-1} \int_t^{\tau_1} \sigma(u) e^{\kappa(\tau_1-t)} \phi'(\theta(u)) du.$$

Proof. We prove first that for a real-valued measurable and bounded function $f(t)$, it holds that (for $t < \tau \leq \tau_{\max} \infty$)

$$(4.6) \quad \mathbb{E}_\theta \left[\int_t^\tau f(u) dL(u) \mid \mathcal{F}_t \right] = \int_t^\tau f(u) \phi'(\theta(u)) du.$$

We calculate using Bayes' Formula, the independent increment property of the Lévy process and the definition of the moment generating function of L :

$$\begin{aligned}
\mathbb{E}_\theta \left[\int_t^\tau f(u) dL(u) \mid \mathcal{F}_t \right] &= \mathbb{E} \left[\int_t^\tau f(s) dL(s) \frac{Z^\theta(\tau)}{Z^\theta(t)} \right] \\
&= \exp \left(\int_t^\tau \phi(\theta(u)) du \right) \frac{d}{d\lambda} \mathbb{E}_\theta \left[\exp \left(\int_t^\tau \lambda f(u) + \theta(u) dL(u) \right) \right] \Big|_{\lambda=0} \\
&= \exp \left(\int_t^\tau \phi(\theta(u)) du \right) \frac{d}{d\lambda} \exp \left(\int_t^\tau \phi(\lambda f(u) + \theta(u)) du \right) \Big|_{\lambda=0} \\
&= \int_t^\tau f(u) \phi'(\theta(u)) du.
\end{aligned}$$

Thus, (4.6) holds.

We continue with deriving the futures dynamics: first, split the CAT to obtain

$$\mathbb{E}_\theta \left[\int_{\tau_1}^{\tau_2} T(u) du \mid \mathcal{F}_t \right] = \mathbb{E}_\theta \left[\int_t^{\tau_2} T(u) du \mid \mathcal{F}_t \right] - \mathbb{E}_\theta \left[\int_t^{\tau_1} T(u) du \mid \mathcal{F}_t \right].$$

Inserting the dynamics of the cumulative temperature (2.11) and using the adaptivity of the Lévy process together with (4.6), we find for $t < \tau$

$$\begin{aligned}
\mathbb{E}_\theta \left[\int_t^\tau T(u) du \mid \mathcal{F}_t \right] &= \int_t^\tau s(u) du + \kappa^{-1} (T(0) - s(0)) (e^{\kappa\tau} - e^{\kappa t}) \\
&\quad + \kappa^{-1} \int_0^t \sigma(u) (e^{\kappa(\tau-u)} - e^{\kappa(t-u)}) dL(u) \\
&\quad + \kappa^{-1} \mathbb{E}_\theta \left[\int_t^\tau \sigma(u) (e^{\kappa(\tau-u)} - 1) dL(u) \mid \mathcal{F}_t \right] \\
&= \int_t^\tau s(u) du + \kappa^{-1} (T(0) - s(0)) (e^{\kappa\tau} - e^{\kappa t}) \\
&\quad + \kappa^{-1} \int_0^t \sigma(u) (e^{\kappa(\tau-u)} - e^{\kappa(t-u)}) dL(u) \\
&\quad + \kappa^{-1} \int_t^\tau \sigma(u) (e^{\kappa(\tau-u)} - 1) \phi'(\theta(u)) du
\end{aligned}$$

Appealing to the explicit dynamics (2.10) for the temperature yields the result. \square

A straightforward application of the Itô Formula for semimartingales (see Ikeda and Watanabe [18]) gives that the $dL(t)$ -term in the differential of $F(t, \tau_1, \tau_2)$ is given by

$$(4.7) \quad \Sigma(t, \tau_1, \tau_2) := \kappa^{-1} (e^{\kappa(\tau_2-t)} - e^{\kappa(\tau_1-t)}) \sigma(t).$$

The function $\Sigma(t, \tau_1, \tau_2)$ can be interpreted as the volatility of the futures dynamics. In the case of $L = B$, e.g., a Brownian dynamics, the Esscher transform coincides with a Girsanov transform, and under the risk-neutral probability the futures dynamics will simply become

an additive process without drift given by

$$(4.8) \quad dF(t, \tau_1, \tau_2) = \Sigma(t, \tau_1, \tau_2) dW(t).$$

Here W being a Brownian motion with respect to Q . Using the Gaussian structure of the distribution of F , we can in this case calculate an explicit pricing formula for a call option written on the futures contract having exercise date $\tau \leq \tau_1$ and strike price K .

Proposition 4.2. *If $L = B$, then the price of a call option at time t on the futures contract with exercise date $t \leq \tau \leq \tau_1$ and strike price K is*

$$(4.9) \quad C(t) = e^{-r(\tau-t)} (F(t) - K) \Phi(d) + \frac{\Sigma_{t,\tau}}{\sqrt{2\pi}} e^{-d^2/2},$$

where

$$d = \frac{F(t) - K}{\Sigma_{t,\tau}}, \quad \Sigma_{t,\tau} := \sqrt{\int_t^\tau \Sigma^2(u, \tau_1, \tau_2) du},$$

and Φ the cumulative probability function for the standard normal distribution.

Proof. This is a straightforward calculation using the properties of the normal distribution. \square

We note that the option value $C(t)$ in Prop. 4.2 is independent of the market price of risk θ as expected, since in the case of $L = B$ we can completely hedge the option using the underlying futures contract. If, however, we model the temperature dynamics using a Lévy process, we are led to the calculation of the following expression:

$$(4.10) \quad C(t) = e^{-r(\tau-t)} \mathbb{E}_\theta [\max(F(\tau, \tau_1, \tau_2) - K, 0) | \mathcal{F}_t],$$

which will depend on the specification of θ . In this case, the Lévy process introduces an incompleteness of the market preventing the possibility to hedge the option. Thus, we have many arbitrage-free prices, which we have parametrized by θ . This market price of risk can be fixed from fitting the theoretical futures prices in Thm. 4.1 to the historical ones (see Benth, Ekeland, Hauge and Nielsen [5] for similar considerations in the energy market). The expectation in (4.10) can next be calculated using numerical integration or Monte Carlo simulation. In order to do so, it is necessary to know the distributional properties of the random variable $T(\tau)$, or, equivalently, the random variable

$$Y(\tau) := \int_0^\tau \sigma(u) e^{\kappa(\tau-u)} dL(u)$$

with respect to Q^θ . From Benth and Saltyte-Benth [6, Cor. 4.4], we have that the characteristic function Ψ of $Y(\tau)$ is given by $\mathbb{E}_\theta [\exp(i\lambda Y(\tau))] = \exp(\Psi(\lambda))$, where

$$\Psi(\lambda) = \int_0^\tau \psi(\lambda \sigma(t) e^{\kappa(\tau-t)} - i\theta(t)) dt - \int_0^\tau \psi(-i\theta(t)) dt,$$

and ψ is the characteristic function of $L(1)$ and $i = \sqrt{-1}$. Note that for the case of the generalized hyperbolic distribution we know the moment generating function ϕ , and thus also the characteristic function $\psi(\lambda) = \phi(i\lambda)$. We can now find the distribution for a

given model (for instance the generalized hyperbolic class) by numerical inversion of the characteristic function Ψ .

5. CONCLUSIONS AND TOPICS FOR FURTHER RESEARCH

We have suggested a non-Gaussian Ornstein-Uhlenbeck model (2.5) to describe the stochastic dynamics of temperature. The model includes a seasonal volatility and a Lévy process generating the innovations, slightly extending other Ornstein-Uhlenbeck models proposed in the literature for describing temperature dynamics (Dornier and Querel [16], Alaton *et al.* [3] and Brody *et al.* [12]), but much less parsimonious than the autoregressive ARCH model of Campbell and Diebold [13].

An empirical study of Norwegian temperature data confirmed the mean-reversion properties of (deseasonalized) daily temperature data. After taking the squares of the residuals, we observed that the correlogram had a wavy form clearly suggesting a seasonally varying variance. We modelled this by an empirically based function. The obtained residuals after dividing by the square-root of this variance function were not normally distributed for some of the measurement stations. We proposed to use the generalized hyperbolic distributions, which explained the skewness and (semi-)heavy tails in the empirical distributions rather well. Introducing the generalized hyperbolic class led us to a Lévy process dynamics in the stochastic model for temperature. We showed that the Norwegian temperature data have fractional characteristics, but argued in favour of a model which is not based on an Ornstein-Uhlenbeck process driven by a fractional Brownian motion.

We calculated an explicit arbitrage-free dynamics for futures prices based on cumulative average temperature using our proposed model. The futures price dynamics will be governed by the market price of risk, which has to be estimated from historical futures prices. Letting the Lévy process in the temperature dynamics be Brownian motion, we derived an arbitrage-free price for a call option written on such futures. The general case must be treated by numerical methods, and we provided a representation of the characteristic function of the temperature dynamics under the risk-neutral probability, a key ingredient to find the density function necessary for pricing.

The empirical studies revealed that the residual noise had positive correlation at lag 1, and negative correlation at lag 2, both being significant (see e.g. Subsection 3.4). Our model was not able to explain this fact, which was observed for all measurement locations. A moving-average time series model could, for instance, be used to capture this effect, however, the corresponding continuous-time dynamics would become a delayed stochastic differential equation. Furthermore, for the squares of these residuals, one could observe a decaying autocorrelation function for small lags (see Subsection 3.4). Following Bollerslev [11], one could try out a GARCH-model to capture this behaviour.

From a risk management point of view, it is of importance to have a multidimensional model for temperature covering all locations of interest. One should here investigate different (possibly time dependent) models for a multidimensional residual model. From our empirical studies we have an understanding of the marginal behaviour, but it is a challenge to model the correlation structure and investigate possible time-dependencies in this.

REFERENCES

- [1] Abramowitz, M. and Stegun, I. A. (1972). *Handbook of Mathematical Functions*, Dover Publications, New York.
- [2] Aivazian, S. A. and Mkhitarian, V. S. (1998). *Applied Statistics and Essentials of Econometrics*. UNITY, Moscow (in Russian).
- [3] Alaton, P., Djehiche, B. and Stillberger, D. (2002). On modelling and pricing weather derivatives. *Appl. Math. Finance*, **9**(1), 1–20.
- [4] Benth, F. E. (2003). On arbitrage-free pricing of weather derivatives based on fractional Brownian motion. *Appl. Math. Finance*, **10**(4), 303–324.
- [5] Benth, F. E., Ekeland, L., Hauge, R. and Nielsen, B. F. (2003). A note on arbitrage-free pricing of forward contracts in energy markets. *Appl. Math. Finance*, **10**(4), 325–336.
- [6] Benth, F. E. and Saltyte-Benth, J. (2003). The normal inverse Gaussian distribution and spot price modelling in energy markets. To appear in *Intern. J. Theor. Appl. Finance*.
- [7] Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *J. Political Econom.*, **81**, 637–654.
- [8] Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type. *Finance & Stochastics*, **2**, 41–68.
- [9] Barndorff-Nielsen, O. E. and Shephard, N. (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **63**(2), 167–241.
- [10] Barndorff-Nielsen, O. E. and Shephard, N. (2003). *Financial volatility: Stochastic volatility and Lévy based models*. Cambridge University Press, Cambridge. Monograph in preparation.
- [11] Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *J. Econometrics*, **31**, 307–327.
- [12] Brody, D. C., Syroka, J. and Zervos, M. (2002). Dynamical pricing of weather derivatives. *Quantit. Finance*, **3**, 189–198.
- [13] Campbell, S. D. and Diebold, F. X. (2002). Weather forecasting for weather derivatives. *Manuscript*. Available on <http://www.ssc.upenn.edu/fdiebold/papers/papers.html>. Version: December 4, 2002.
- [14] Davis, M. H. A. (2001). Pricing weather derivatives by marginal value. *Quantit. Finance*, **1**, 305–308.
- [15] Doornik, J. A. (1998). *Ox: Object Oriented Matrix Programming 2.0*. Timberlake, London.
- [16] Dornier, F. and Querel, M. (2000). Caution to the wind. *Energy Power Risk Manag.*, Weather risk special report, August, 30–32.
- [17] Eberlein, E. and Keller, U. (1995). Hyperbolic distributions in finance. *Bernoulli*, **1**, 281–299.
- [18] Ikeda, N. and Watanabe, S. (1981). *Stochastic Differential Equations and Diffusion Processes*, North-Holland and Kodansha, Tokyo.
- [19] Karatzas, I., and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*, Springer-Verlag, New York.
- [20] Merton, R. C. (1973). The theory of rational option pricing. *Bell J. Econom. Manag. Sci.*, **4**, 141–183.
- [21] Rydberg, T. H. (1997). The normal inverse Gaussian Lévy process: simulation and approximation. *Commun. Statist.-Stochastic Models*, **13**(4), 887–910.
- [22] Shiryaev, A.N. (2000). *Essentials of Stochastic Finance: Facts, Models, Theory*, World Scientific, Singapore.
- [23] Syroka, J. I. and Toumi, R. (2001). Scaling and persistence in observed and modelled surface temperature. *Geophys. Res. Lett.* **28**, 3255–3259.

[Fred Espen Benth]
Centre of Mathematics for Applications
University of Oslo
P.O. Box 1053, Blindern
N-0316 Oslo, Norway
and
Agder University College
Department of Economics and Business Administration
Serviceboks 422
N-4604 Kristiansand, Norway
e-mail: fredb@math.uio.no
URL: <http://www.math.uio.no/~fredb/>

[Jūratė Šaltytė-Benth]
Centre of Mathematics for Applications
University of Oslo
P.O. Box 1053, Blindern
N-0316 Oslo, Norway
and
Department of System Research
Klaipėda University
H. Manto 84
LT-5808 Klaipėda, Lithuania
e-mail: jurate@math.uio.no

APPENDIX A. FIGURES




FIGURE 1. Average daily temperatures with the seasonal component for Alta and Bergen




FIGURE 2. Scatter plot from regressing today's deseasonalized temperature against yesterday's deseasonalized temperature (left), and today's absolute temperature against yesterday's deseasonalized (right)




FIGURE 3. The residuals and squared-residuals after linear regression of deseasonalized daily temperatures. The plots show Alta and Bergen




FIGURE 4. Empirical density (broken line) and fitted normal (complete line) for the residuals after regression. The plots display Alta and Bergen, where a logarithmic scale for the frequencies are used in the second and fourth plot. The empirical densities are plotted using a Gaussian kernel smoother

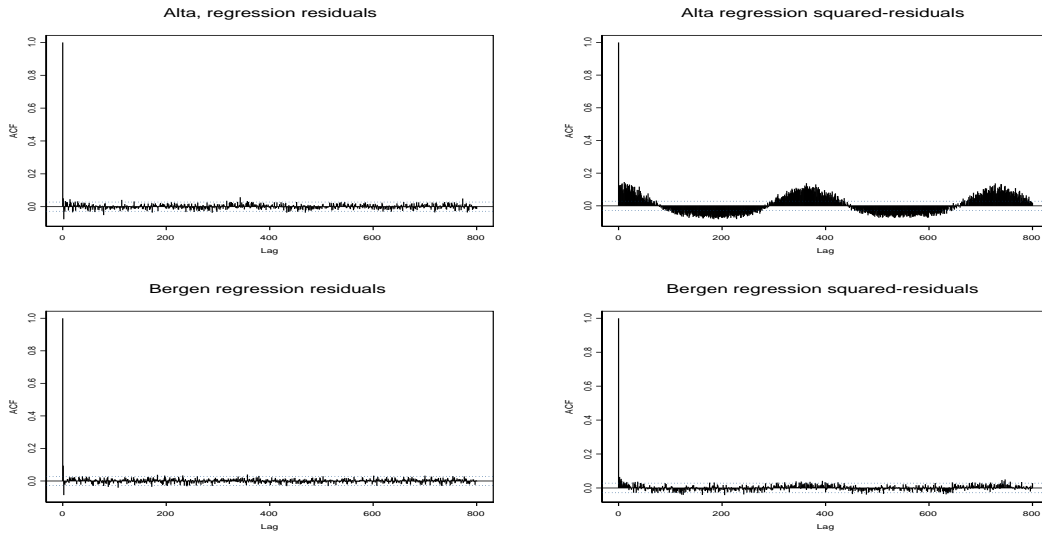


FIGURE 5. The empirical autocorrelation function for the residuals and squared residuals after regression for Alta and Bergen




FIGURE 6. The empirical and smoothed seasonal variability function. The plots display Alta and Bergen




FIGURE 7. The residuals and squared-residuals after dividing out the seasonal variation component. The plots display Alta and Bergen

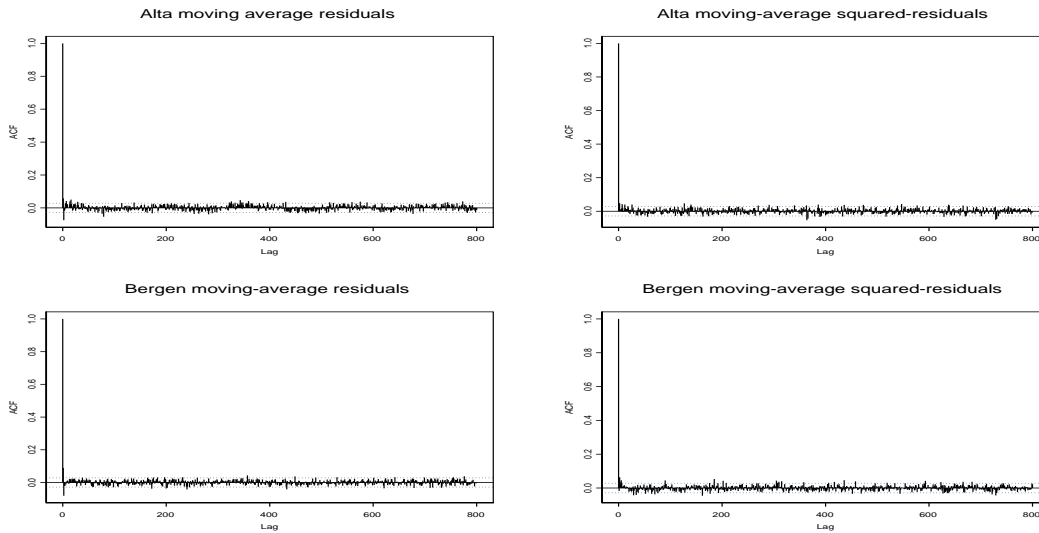


FIGURE 8. The empirical autocorrelation function for the residuals and squared residuals after dividing by the seasonal variation σ_t for Alta and Bergen




FIGURE 9. Empirical density (broken line) and fitted normal distribution (complete line) for the residuals after dividing by the seasonality variation. The plots display Alta and Bergen, where a logarithmic scale for the frequencies are used in the second and fourth plot. The empirical densities are plotted using a Gaussian kernel smoother




FIGURE 10. Empirical density (broken line) and fitted generalized hyperbolic distribution (complete line) for the residuals after dividing by the seasonality variation. The plots display Alta and Bergen, where a logarithmic scale for the frequencies are used in the second and fourth plot. The empirical densities are plotted using a Gaussian kernel smoother




FIGURE 11. Empirical density (broken line) and fitted generalized hyperbolic distribution (complete line) for the residuals after dividing by the seasonality variation. The plots display Kristiansand and Oslo, where a logarithmic scale for the frequencies are used in the second and fourth plot. The empirical densities are plotted using a Gaussian kernel smoother