Energy Economics 34 (2012) 592–602

Contents lists available at SciVerse ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneco

A critical view on temperature modelling for application in weather derivatives markets

Jūratė Šaltytė Benth [a,b,*], Fred Espen Benth [c,d]

[a] Institute of Clinical Medicine, Campus Ahus, University of Oslo, Norway
[b] HØKH, Research Centre, Akershus University Hospital, Norway
[c] Centre of Mathematics for Applications, University of Oslo, Norway
[d] School of Management, University of Agder, Norway

ARTICLE INFO

Article history:
Received 12 August 2010
Received in revised form 19 September 2011
Accepted 25 September 2011
Available online 8 October 2011

JEL classification:
C18
C51

Keywords:
Temperature
Time series model
Weather derivatives
Seasonality
GARCH

ABSTRACT

In this paper we present a stochastic model for daily average temperature. The model contains seasonality, a low-order autoregressive component and a variance describing the heteroskedastic residuals. The model is estimated on daily average temperature records from Stockholm (Sweden). By comparing the proposed model with the popular model of Campbell and Diebold (2005), we point out some important issues to be addressed when modelling the temperature for application in weather derivatives market.

1. Introduction

In recent years there has been a growing interest in modelling the dynamics of surface air temperature with application in pricing weather derivatives. We follow up this stream of research with a critical discussion on model building and estimation, contrasting two stochastic models proposed by Campbell and Diebold (2005) and Benth and Šaltytė-Benth (2007). Both models are based on a seasonal autoregressive (AR) process, but with significant differences in structure which influences their applicability in relation to temperature derivatives. The two models are widely used in the field, and are similar to or nest a number of related models, see, for example, Dornier and Querel (2000), Alaton et al. (2002), Cao and Wei (2004) to mention a few. The performance (in terms of forecasting weather indices) of various models for temperature dynamics, including the two considered here, was compared by Oetoma and Stevenson (2005), Svec and Stevenson (2007), Papazian and Skiadopoulos (2010), Zapranis and Alexandridis (2008), Schiller et al. (forthcoming). Our main goal is to point out the principle differences between the models of Campbell and Diebold, (2005) and Benth and Šaltytė-Benth (2007).

* Corresponding author at: Mail Drawer 95, HØKH, Akershus University Hospital, 1478 Lørenskog, Norway. Tel.: +47 48224431.
E-mail address: jurates@medisin.uio.no (J. Šaltytė Benth).

0140-9883/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.eneco.2011.09.012

At the Chicago Mercentile Exchange (CME) there is an organized trade in weather futures and options. In particular, the CME offers trade in futures contracts written on temperature indices measured at various locations world wide, providing financial instruments to hedge weather risk exposure. The locations are major cities in the US, Canada, Europe and Asia. The temperature indices measure the daily cumulative average temperature (CAT), the cumulative heating-degree days (HDD) or the cumulative cooling-degree days (CDD). The reference temperature is taken as the average of the daily minimum and maximum temperature, which we further refer to as temperature.

More specifically, if we denote the temperature on day t by $Z(t)$, then the CAT index over a measurement period $[T_1, T_2]$ is defined as

$$\mathrm{CAT}(T_1,T_2)=\sum_{t=T_1}^{T_2} Z(t), \tag{1}$$

where the measurement period is typically a given month or season. At CME, CAT futures are traded on European cities for measurement periods in warm season. The HDD index measures the demand for heating in the cold period of the year, and is defined as the cumulative amount of average temperatures below a threshold over a measurement period. That is, one aggregates $max(c-Z(t),\ 0)$, where the threshold c is 65°F or 18 °C. The CDD index analogously aggregates $max(Z(t)-c, 0)$ and measures the demand for air-conditioning cooling. The CDD and HDD futures are traded for US cities.

The temperature futures contracts are financially settled proportionally to the underlying index at the end of the measurement period. To assess the risk in trading such contracts and to be able to settle reasonable futures prices, one needs precise models for the temperature dynamics. A model should incorporate such properties as possible trend due to global warming and/or urbanisation, seasonal component describing periodic temperature variations related to cold and warm seasons, AR properties for temperature changes, and seasonal variations in residuals. In this study, we follow up the analysis from papers Benth and Šaltytė-Benth (2007), Benth et al. (2007), Šaltytė et al. (2007), and contrast it to the approach of Campbell and Diebold (2005).

As observed by Campbell and Diebold (2005) for US temperature data, and later confirmed for Swedish, Norwegian and Lithuanian temperatures (see above-cited papers), there is a clear seasonality in the temperature variations after removing trend, seasonal variations and AR effects from the data. The model for temperature proposed by Benth and Šaltytė Benth contains (linear) trend, seasonal component, low-order AR process and a seasonal variance component describing the remaining heteroskedasticity in temperature residuals. It differs from the model of Campbell and Diebold (2005) first of all in its simplicity. Beside trend and seasonal component, Campbell and Diebold propose to use a high-order AR time series model with seasonal generalized AR conditional heteroskedesticity (GARCH) model for the residuals. We are able to show that the model by Benth and Šaltytė Benth is sophisticated enough to explain the basic stylized facts of temperatures just as well as the parameter-intensive alternative proposed by Campbell and Diebold (2005).

The contribution of this paper is threefold. First, we critically review the process of modelling the temperatures. In this paper we promote a stepwise procedure used in Benth and Šaltytė-Benth (2007), Benth et al. (2007) and Šaltytė et al. (2007), where one models and estimates each component step-by-step. This turns out to be advantageous in order to build a confident model explaining the various stylized facts of temperature. In particular, such an approach leads to a very low-order AR structure in the temperature dynamics, in contrast to the approach of Campbell and Diebold (2005) suggesting to use an AR model with 25 lags. We argue that one can model temperature dynamics equally good using AR(3) and show that a simpler model explains the temperature evolution very well. A bigger empirical analysis also shows that the model by Benth and Šaltytė Benth explains extremely good the historical index values of CAT, HDD and CDD.

As a second contribution, we discuss the major role played by the mean of temperature in the context of weather derivatives. The main factor explaining the indices CAT, HDD and CDD turns out to be the seasonal mean temperature, as we demonstrate later on. This is not surprising, taking into consideration the relatively strong mean reversion of temperatures along with indices averaging temperatures over intervals like months. In the model proposed by Benth and Šaltytė Benth, the seasonal mean is modelled explicitly, and is directly estimated fromtemperature observations. In this way one obtains a confident model for the seasonality of temperature. Campbell and Diebold (2005) choose to model it *indirectly*, estimated together with all the other parameters in the model. In the model by Benth and Šaltytė Benth, one regresses the deseasonalized temperatures on *deseasonalized* temperatures, that is, the AR structure is modelled after removing the seasonal mean. Campbell and Diebold (2005) choose to regress today's *deseasonalized* temperature on the temperature in previous days. Their seasonal function will then *not* be the seasonal mean, but merely a seasonal component. We demonstrate how one can compute the seasonal mean from the model of Campbell and Diebold, involving the AR parameters and thus leading to potentially increased uncertainty in parameter estimates.

The third contribution of the paper is a multiplicative seasonal stochastic volatility model. Instead of using an additive GARCH process in modelling the seasonal heteroskedastic residuals as Campbell and Diebold (2005) do, we suggest using a product between a seasonal deterministic function and a classical GARCH process instead. With a multiplicative structure one avoids potential problems related to the positivity of variance. Moreover, no new estimation procedure is required to estimate the GARCH component, leading to a model which is simpler to fit and therefore more practically applicable.

The paper is organized as follows. First, we state the model for the daily temperature variations and discuss in detail the different components of it. Then we describe the data and estimate the model. Next, we validate the proposed model on out-of-sample data and apply it for forecasting different temperature indices. We end our paper by discussing continuous-time models and weather derivatives pricing.

2. A general model for temperature dynamics

We present a general time series model for the temperature dynamics, which is nesting many of the existing models. For modelling of temperature, we suggest to use a time series decomposition approach, where the time series is decomposed into different components like trend, seasonality, an AR process (so-called cyclic component) and residual term, all appearing in observed data simultaneously. By estimating and eliminating different components of time series step-by-step and examining all intermediate residuals, one gets a good insight into the data structure and is likely able to come up with a precise model.

We consider the following model for temperature (see Benth et al. (2008)):

$$Z(t) = \mu(t) + \varepsilon(t), \tag{2}$$

where $\mu(t)$ and $\varepsilon(t)$ denote, respectively, the mean and the residual process at time $t = 1, \ldots, \tau$. Here

$$\mu(t) = S(t) + \sum_{i=1}^{p} \alpha_i (Z(t-i) - S(t-i)), \tag{3}$$

where $S(t)$ is a deterministic function and α_i, $i = 1, \ldots, p$, are the parameters of the AR(p) process. A more general autoregressive moving average process (ARMA) could be considered instead of AR, but the empirical analysis suggests that there is no need for such an extension. The AR parameters can in general be time-dependent. The stability analysis of AR(1) process was performed in Benth et al. (2007) for Stockholm temperatures. There were no significant differences observed among the regression parameters estimated for different years or seasons. We therefore assume that the mean reversion for Stockholm temperatures is stable over time.

The deterministic function $S(t)$ plays the role of the long-term average of the temperature, towards which the temperature mean reverts due to the AR structure. One could think of fitting ARMA process directly on the temperature observations. However, Oetoma and Stevenson (2005) have shown that a conventional ARMA model without controlling for long-term trend and seasonality does not outperform alternative models.

Another way to represent Eq. (2) is to write

$$Z(t) - S(t) = \sum_{i=1}^{p} \alpha_i (Z(t-i) - S(t-i)) + \varepsilon(t),$$

where it is assumed that the *deseasonalized* temperature follows an AR(p) process, i.e. todays' deseasonalized temperature is regressed on the p previous days' deseasonalized temperatures. As long as the residual process