Energy Economics 55 (2016) 112–126

Contents lists available at ScienceDirect

Energy Economics

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

A consistent two-factor model for pricing temperature derivatives

Andreas Groll [a,*], Brenda López-Cabrera [b], Thilo Meyer-Brandis [a]

[a] Mathematical Institute, University of Munich, Theresienstrasse 39, 80333 Munich, Germany
[b] Ladislaus von Bortkiewicz Chair of Statistics, Humboldt-Universität zu Berlin, Spandauer Straße 1, 10178 Berlin, Germany

ARTICLE INFO

Article history:
Received 1 June 2015
Received in revised form 15 December 2015
Accepted 30 December 2015
Available online 26 January 2016

JEL classification:
C490
C130
G19
G29
G22
Q59

Keywords:
Factor models
Consistency
Pricing and hedging
Weather derivatives
Market price of risk

ABSTRACT

In the past decade, the Chicago Mercantile Exchange began to trade weather derivatives to hedge weather risk. The pricing of weather derivatives is challenging since the underlying is not tradable and thus classical arbitrage approaches have to be used with caution. In typical pricing approaches all information available to the market is assumed to be incorporated in the underlying and thus forward-looking information about non-tradable assets such as meteorological forecasts is often ignored. In this article, we analyze a new pricing methodology for temperature derivatives that accounts for forward-looking information. More precisely, we provide an empirical back-up for the theoretical framework of so-called consistent factor models for temperature forecast curves introduced previously in the literature and put this pricing approach into practice. First, we perform a thorough statistical analysis of meteorological forecast curve data. Second, based on this analysis we propose a specific consistent two-factor model, derive explicit temperature derivative prices, and calibrate the market price of risk (MPR). The power of the model is demonstrated against alternative pricing models. This confirms that at least parts of the irregularity of the MPR observed in earlier studies are not due to irregular risk perception but rather due to information misspecification.

1. Introduction

In the last years weather derivatives (WD) have emerged to hedge weather variability. This leads to the question how such derivatives are priced and hedged. In contrast to other assets, the pricing of weather derivatives has some challenges since the underlying, contingent on temperature or rain, is not tradeable and the classical Black, Merton and Scholes framework fails since hedging principles cannot be applied. Different streams for pricing weather derivatives are found in the literature: econometric modeling of the underlying dynamics (Campbell and Diebold, 2005), risk neutral pricing inspired by methodologies from financial markets (Benth et al., 2007), equilibrium models (Cao and Wei, 2004), indifference pricing models (Brockett et al., 2006), index modeling and burn analysis (Jewson and Brix, 2005), to mention some of them.

In the context of no-arbitrage pricing, there has been some work to calibrate and study the complex structure of the risk premium, or equivalently the market price of risk (MPR), in temperature derivative prices, see for example the studies of Cao and Wei (2004), Huang-Hsi et al. (2008) and Richards et al. (2004)). While most of the articles on temperature derivative pricing assume a zero or constant MPR (Dorfleitner and Wimmer, 2010; Cao and Wei, 2004; Huang-Hsi et al., 2008; Richards et al., 2004; Alaton et al., 2002), a more differentiated analysis conducted in Härdle and López-Cabrera (2012), Benth et al. (2008), Benth and Benth (2011, 2012) and Härdle et al. (2011) actually reveals a complex time varying and stochastic behavior of the MPR. However, we believe that at least parts of the irregular behavior of the MPR are not due to irregular risk pricing of the market but due to an information misspecification of most models that are used to calibrate market prices: forward looking information about the temperature available to the market is not taken into account in the information modeling. The usual assumption that all information available to the market is incorporated in the past behavior of the underlying, i.e. the information filtration is generated by the underlying, might be acceptable for storable assets (classical financial markets). However, for non-storable underlyings (like temperature or electricity) this assumption is fundamentally wrong: a substantial amount of forward looking information available to the market, like meteorological weather forecast, is

* Corresponding author. Tel.: +49 89 2180 4697; fax: +49 89 2180 4452.
E-mail addresses: groll@math.lmu.de (A. Groll), lopezcab@wiwi.hu-berlin.de (B. López-Cabrera), meyerbra@math.lmu.de (T. Meyer-Brandis).
URLs: http://www.fm.mathematik.uni-muenchen.de/personen/phd_postdoc/groll/index.html (A. Groll), https://www.wiwi.hu-berlin.de/professuren/quantitativ/statistik/members/personalpages/bl (B. López-Cabrera), http://www.fm.mathematik.uni-muenchen.de/personen/professors/meyer_brandis/index.html (T. Meyer-Brandis).

http://dx.doi.org/10.1016/j.eneco.2015.12.020
0140-9883/

not reflected in the past evolution of temperature. Even though meteorological forecasts usually have a short time horizon and thus improve pricing of weather derivatives especially when those are close to or in the delivery period, an appropriate model for the pricing of temperature derivatives should take into account this information, which is available to the market participants.

There are few studies dealing with the incorporation of meteorological forecast into weather derivative pricing. Alaton et al. (2002) suggests, without a model, to incorporate forecasts for short-term pricing. Jewson and Brix (2005) describe how to use single and ensemble forecasts to derive probabilistic weather forecasts for WD pricing. Yoo (2003) incorporates the seasonal forecast in the temperature process, by assuming the unconditional mean temperature as a linear combination of above-normal (warm), near-normal, and below-normal (cool) mean temperature processes. Benth and Meyer-Brandis (2009) apply the theory of enlargement of filtration to describe all information available in the market and with it estimate information premiums. The studies of Ritter et al. (2011) and Härdle et al. (2012) applied the model from Benth and Meyer-Brandis (2009) and show that incorporating weather forecast gives better fittings of market prices. Another empirical study of the information premium in the context of electricity markets is described in Benth et al. (2013), where the existence and the size of the corresponding information premium are analyzed. However, the challenge of working with the theory of enlargement of filtrations is the analytic tractability of the forward looking information. In Dorfleitner and Wimmer (2010) an index modeling approach is used and it is also shown that weather forecasts significantly influence prices.

In this manuscript we follow-up on the pricing approach for temperature derivatives outlined in Hell et al. (2012). There the idea is to model the dynamics of complete temperature forecast curves by so-called consistent factor models. The temperature dynamics is given by the short end of the curve, which then accounts for the information about meteorological forecasts available to the market participants when pricing temperature derivatives. More precisely, given a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_{t\geq 0}, \mathbb{P})$ fulfilling the usual conditions, it is assumed that the meteorological forecast $f(t;T)$ at time t of the temperature $\tau(T)$ at time T given all available information $\mathcal{F}_t$ is an unbiased estimator of the temperature in the sense that

$$f(t;T) = \mathrm{E}[\tau(T)|\mathcal{F}_t]. \tag{1}$$

For fixed forecast time T, the forecast process $f(\cdot;T)$ is thus a martingale under the real world probability measure $\mathbb{P}$ with respect to the flow of available information represented by $\mathcal{F}_{t\geq 0}$. Next, the stochastic evolution of the forecast curves is modeled by a factor model

$$f(t,T) = H(T-t, Z(t)),$$

where

- $H(x,z): \mathbb{R}_+ \times \mathbb{R}^m \to \mathbb{R}$ is a given curve family;
- $x = T - t$ is time to forecast time;
- $Z(t)$ is an $\mathbb{R}^m$-valued factor process given by an Itô diffusion

$$dZ(t) = b(Z(t))\ dt + \sigma(Z(t))dW(t), \quad Z(0) = z_0,$$

with $W(t)$ a d-dimensional $\mathcal{F}_t$-Brownian motion.

By setting $T = t$, we then obtain the dynamic model for the temperature $\tau(t) = f(t;t)$ which now includes driving factors from the complete forecast curve.

A factor model is called consistent if $H(T-t, Z(t))$ is an $\mathcal{F}_t$-martingale for any fixed $T \geq 0$ in accordance with (1). Given a curve family $H(x,z)$, the consistency requirement imposes certain restrictions on admissible drift coefficients $b(z)$ and volatility coefficients $\sigma(z)$ which, for certain types of families $H(x,z)$, are characterized in Hell et al. (2012).

While in Hell et al. (2012) the general theoretical framework for the pricing of temperature derivatives by consistent factor models is established, the purpose of this manuscript is to give an empirical back-up and the implementation into practice of this approach. To this end, the first step will be a thorough descriptive statistical analysis of temperature forecast curve time series and the identification of stylized facts such as seasonality, seasonal variance, exponential decaying autocorrelation, stationarity and mean-reversion. Further, a functional principal component analysis (FPCA) is employed in order to understand the structure of the driving factors of forecast curve dynamics such as mean-reversion in the long horizon and tilting or bending shapes in the short horizon. In the second part of the article, motivated by the statistical investigations, we then analyze in detail a specific consistent two-factor model, which is based on a generalization of the Nelson–Siegel curve family. We confirm that this model, introduced in Hell et al. (2012) as an example, indeed captures well the stylized facts of forecast curves while at the same time yields high analytical tractability when computing derivative prices. In particular, the implied model for the temperature extends other models proposed in the WD literature by including a mean-reversion to a stochastic mean level that accounts for forecast information. After specifying and implementing an iterative two-step algorithm for the estimation of our model, we derive explicit prices for temperature derivatives written on Cumulative Average Temperature (CAT), Heating Degree Days (HDD) and Cooling Degree Days (CDD). Finally, we calibrate the market price of risk (MPR) of temperature derivatives (CAT, CDD, HDD) traded at the Chicago Mercantile Exchange (CME). The results show that our forecast curve model clearly outperforms the popular Ornstein-Uhlenbeck (OU) model that ignores forward looking information in reproducing in- and out-of-sample traded futures prices. This confirms that at least parts of the irregularity of the MPR in other studies are due to information misspecification.

Our article is structured as follows. We start by the statistical analysis of temperature and meteorological forecast data in Section 2 to lay the basis for an appropriate model choice. Motivated by the empirical analysis in Section 2, we then analyze a specific consistent two factor model in Section 3 and present an estimation algorithm. The pricing of temperature derivatives and the calibration of the MPR of the proposed factor model, as well as the comparison with an OU model, are developed in Section 4. Section 5 concludes the article.

All computations in this article were carried out in the statistical program R (R Core Team, 2013). The temperature and weather derivative data were obtained from Bloomberg and are available in the Risk Data Center of the CRC 649 Economic Risk (http://sfb649.wiwi.huberlin.de/). WeatherOnline provides the meteorological forecast data. Further, to simplify notation, in the following dates are denoted in the yyyymmdd format.

2. Empirical analysis of temperature forecast curves

2.1. Data

In this section we present an empirical analysis of the temperature and point temperature forecast data of the cities New York and Berlin. The temperature data used in this study for Berlin-Tempelhof airport and New York-Laguardia airport are the daily average temperatures from 19480101 - 20111231 and 19700101 - 20121231 (yyyymmdd), respectively, and are provided by the CME website.

Meteorological forecast data are derived from WeatherOnline.[1] These data consist of point forecasts of the minimal and maximal temperatures for Berlin-Tempelhof airport and New York-Laguardia airport from 1 to 14 days in advance (14 days), calculated every day for the time period 20081229 - 2010211 and 20081229 - 201301003, respectively.

[1] We thank Dr. Ulrich Römer and Herrad Werner for providing us the data. Note here that as our forecast and temperature data had to be obtained from different providers, this bears a risk of inconsistencies in the data.